arXiv:1609.09449v1 [cs.SY] 29 Sep 2016

A Cross Entropy based Stochastic Approximation Algorithm for Reinforcement Learning with Linear Function Approximation

Ajin George Joseph
Department of Computer Science and Automation, Indian Institute of Science, Bangalore, India, ajin@csa.iisc.ernet.in

Shalabh Bhatnagar
Department of Computer Science and Automation, Indian Institute of Science, Bangalore, India, shalabh@csa.iisc.ernet.in

In this paper, we provide a new algorithm for the problem of prediction in Reinforcement Learning, *i.e.*, estimating the Value Function of a Markov Reward Process (MRP) using the linear function approximation architecture, with memory and computation costs scaling quadratically in the size of the feature set. The algorithm is a multi-timescale variant of the very popular Cross Entropy (CE) method which is a model based search method to find the global optimum of a real-valued function. This is the first time a model based search method is used for the prediction problem. The application of CE to a stochastic setting is a completely unexplored domain. A proof of convergence using the ODE method is provided. The theoretical results are supplemented with experimental comparisons. The algorithm achieves good performance fairly consistently on many RL benchmark problems. This demonstrates the competitiveness of our algorithm against least squares and other state-of-the-art algorithms in terms of computational efficiency, accuracy and stability.

Key words: Reinforcement Learning, Cross Entropy, Markov Reward Process, Stochastic Approximation, ODE Method, Mean Square Projected Bellman Error, Off Policy Prediction

1. Introduction and Preliminaries

In this paper, we follow the Reinforcement Learning (RL) framework as described in [1, 2, 3]. The basic structure in this setting is the discrete time Markov Decision Process (MDP) which is a 4-tuple ($\mathbb{S}$, $\mathbb{A}$, R, P), where $\mathbb{S}$ denotes the set of *states* and $\mathbb{A}$ is the set of *actions*. $\mathrm{R} : \mathbb{S} \times \mathbb{A} \times \mathbb{S} \to \mathbb{R}$ is the *reward function* where $\mathrm{R}(s, a, s')$ represents the reward obtained in state s after taking action a and transitioning to s'. Without loss of generality, we assume that the reward function is bounded, *i.e.*, $|\mathrm{R}(.,.,.)| \leq \mathrm{R}_{\max} < \infty$. $\mathrm{P} : \mathbb{S} \times \mathbb{A} \times \mathbb{S} \to [0, 1]$ is the *transition probability kernel*, where $\mathrm{P}(s, a, s') = \mathbb{P}(s'|s, a)$ is the probability of next state being s' conditioned on the fact that the current state is s and action taken is a. We assume that the state and action spaces are finite with $|\mathbb{S}| = n$ and $|\mathbb{A}| = b$. A *stationary policy* $\pi : \mathbb{S} \to \mathbb{A}$ is a function

from states to actions, where $\pi(s)$ is the action taken in state s. A given policy π along with the transition kernel P determines the state dynamics of the system. For a given policy π, the system behaves as a Markov Reward Process (MRP) with transition matrix $\mathrm{P}^\pi(s,s')$ = $\mathrm{P}(s,\pi(s),s')$. The policy can also be stochastic in order to incorporate exploration. In that case, for a given $s \in \mathbb{S}$, $\pi(.|s)$ is a probability distribution over the action space $\mathbb{A}$.

For a given policy π, the system evolves at each discrete time step and this process can be captured as a sequence of triplets $(\mathbf{s}_t, \mathbf{r}_t, \mathbf{s}'_t), t \geq 0$, where $\mathbf{s}_t$ is the random variable which represents the current state at time t, $\mathbf{s}'_t$ is the transitioned state from $\mathbf{s}_t$ and $\mathbf{r}_t = \mathrm{R}(\mathbf{s}_t, \pi(\mathbf{s}_t), \mathbf{s}'_t)$ is the reward associated with the transition. In this paper, we are concerned with the problem of *prediction*, *i.e.*, estimating the long run γ-discounted cost $V^\pi \in \mathbb{R}^{\mathbb{S}}$ (also referred to as the *Value function*) corresponding to the given policy π. Here, given $s \in \mathbb{S}$, we let

$$V^\pi(s) \triangleq \mathbb{E}\left[\sum_{t=0}^{\infty} \gamma^t \mathbf{r}_t \Big| \mathbf{s}_0 = s\right], \tag{1}$$

where $\gamma \in [0,1)$ is a constant called the *discount factor* and $\mathbb{E}[\cdot]$ is the expectation over sample trajectories of states obtained in turn from P^π when starting from the initial state s. V^π is represented as a vector in $\mathbb{R}^{|\mathbb{S}|}$. V^π satisfies the well known *Bellman equation* under policy π, given by

$$V^\pi = \mathrm{R}^\pi + \gamma \mathrm{P}^\pi V^\pi \triangleq T^\pi V^\pi, \tag{2}$$

where $\mathrm{R}^\pi \triangleq (\mathrm{R}^\pi(s), s \in \mathbb{S})^\top$ with $\mathrm{R}^\pi(s)$ = $\mathbb{E}\left[\mathbf{r}_t | \mathbf{s}_t = s\right]$, $V^\pi \triangleq (V^\pi(s), s \in \mathbb{S})^\top$ and $T^\pi V^\pi \triangleq ((T^\pi V^\pi)(s), s \in \mathbb{S})^\top$, respectively. Here T^π is called the *Bellman operator*. If the model information, *i.e.*, P^π and R^π are available, then we can obtain the value function V^π by solving analytically the linear system $V^\pi = (I - \gamma \mathrm{P}^\pi)^{-1} \mathrm{R}^\pi$.

However, in this paper, we follow the usual RL framework, where we assume that the model, is inaccessible; only a sample trajectory $\{(\mathbf{s}_t, \mathbf{r}_t, \mathbf{s}'_t)\}_{t=1}^{\infty}$ is available where at each instant t, state $\mathbf{s}_t$ of the triplet $(\mathbf{s}_t, \mathbf{r}_t, \mathbf{s}'_t)$ is sampled using an arbitrary distribution ν over $\mathbb{S}$, while the next state $\mathbf{s}'_t$ is sampled using $\mathrm{P}^\pi(\mathbf{s}_t, .)$ and $\mathbf{r}_t$ is the immediate reward for the transition. The value function V^π has to be estimated here from the given sample trajectory.

To further make the problem more arduous, the number of states n may be large in many practical applications, for example, in games such as chess and backgammon. Such combinatorial blow-ups exemplify the underlying problem with the value function estimation, commonly referred to as the *curse of dimensionality*. In this case, the value function is unrealizable due to both storage and computational limitations. Apparently one has to resort to approximate solution methods where we sacrifice precision for computational tractability. A common approach in this context is the function approximation method [1], where we approximate the value function of unobserved states using available training data.

In the *linear function approximation technique*, a linear architecture consisting of a set of k, n-dimensional feature vectors, $1 \leq k \ll n$, $\{\phi_i \in \mathbb{R}^{\mathbb{S}}\}$, $1 \leq i \leq k$, is chosen *a priori*. For a state $s \in \mathbb{S}$, we define

$$\phi(s) \triangleq \begin{bmatrix} \phi_1(s) \\ \phi_2(s) \\ \vdots \\ \phi_k(s) \end{bmatrix}_{k\times 1}, \quad \Phi \triangleq \begin{bmatrix} \phi(s_1)^\top \\ \phi(s_2)^\top \\ \vdots \\ \phi(s_n)^\top \end{bmatrix}_{n\times k}, \tag{3}$$

where the vector $\phi(\cdot)$ is called the *feature vector*, while Φ is called the *feature matrix*.

Primarily, the task in linear function approximation is to find a weight vector $z \in \mathbb{R}^k$ such that the predicted value function $\Phi z \approx V^\pi$. Given Φ, the best approximation of V^π is its projection on to the subspace $\{\Phi z | z \in \mathbb{R}^k\}$ (column space of Φ) with respect to an arbitrary norm. Typically, one uses the weighted norm $\|.\|_\nu$ where $\nu(\cdot)$ is an arbitrary distribution over $\mathbb{S}$. The norm $\|.\|_\nu$ and its associated linear projection operator Π^ν are defined as

$$\|V\|_\nu^2 = \sum_{i=1}^{n} V(i)^2 \nu(i), \quad \Pi^\nu = \Phi(\Phi^\top D^\nu \Phi)^{-1} \Phi^\top D^\nu, \tag{4}$$

where D^ν is the diagonal matrix with $D^\nu_{ii} = \nu(i), i = 1, \ldots, n$. So a familiar objective in most approximation algorithms is to find a vector $z^* \in \mathbb{R}^k$ such that $\Phi z^* \approx \Pi^\nu V^\pi$.

Also it is important to note that the efficacy of the learning method depends on both the features ϕ_i and the parameter z [4]. Most commonly used features include Radial Basis Functions (RBF), Polynomials,

Fourier Basis Functions [5], Cerebellar Model Articulation Controller (CMAC) [6] *etc*. In this paper, we assume that a carefully chosen set of features is available *a priori*.

The existing algorithms can be broadly classified as (i) Linear methods which include Temporal Difference (TD) [7], Gradient Temporal Difference (GTD [8], GTD2 [9], TDC [9]) and Residual Gradient (RG) [10] schemes, whose computational complexities are linear in k and hence are good for large values of k and (ii) Second order methods which include Least Squares Temporal Difference (LSTD) [11, 12] and Least Squares Policy Evaluation (LSPE) [13] whose computational complexities are quadratic in k and are useful for moderate values of k. Second order methods, albeit computationally expensive, are seen to be more data efficient than others except in the case when trajectories are very small [14].

Eligibility traces [7] can be integrated into most of these algorithms to improve the convergence rate. Eligibility trace is a mechanism to accelerate learning by blending temporal difference methods with Monte Carlo simulation (averaging the values) and weighted using a geometric distribution with parameter $\lambda \in [0, 1)$. The algorithms with eligibility traces are named with (λ) appended, for example TD(λ), LSTD(λ) *etc*. In this paper, we do not consider the treatment of eligibility traces.

Sutton's *TD(λ) algorithm with function approximation* [7] is one of the fundamental algorithms in RL. TD(λ) is an online, incremental algorithm, where at each discrete time t, the weight vectors are adjusted to better approximate the target value function. The simplest case of the one-step TD learning, *i.e.* $\lambda = 0$, starts with an initial vector $\mathbf{z}_0$ and the learning continues at each discrete time instant t where a new prediction vector $\mathbf{z}_{t+1}$ is obtained using the recursion,

$$\mathbf{z}_{t+1} = \mathbf{z}_t + \alpha_{t+1}\delta_t(\mathbf{z}_t)\phi(\mathbf{s}_t).$$

In the above, α_t is the learning rate which satisfies $\sum_t \alpha_t = \infty$, $\sum_t \alpha_t^2 < \infty$ and $\delta_t(z) \triangleq \mathbf{r}_t + \gamma z^\top \phi(\mathbf{s}'_t) - z^\top \phi(\mathbf{s}_t)$ is called the *Temporal Difference (TD)-error*. In on-policy cases where Markov Chain is ergodic and the sampling distribution ν is the stationary distribution of the Markov Chain, then with α_t satisfying the above conditions and with Φ being a full rank matrix, the convergence of TD(0) is guaranteed [15]. But in off-policy cases, *i.e.*, where the sampling distribution ν is not the stationary distribution of the chain, TD(0) is shown to diverge [10].

By applying stochastic approximation theory, the limit point z^*_{TD} of TD(0) is seen to satisfy

$$0 = \mathbb{E}\left[\delta_t(z)\phi(\mathbf{s}_t)\right] = Az - b, \tag{5}$$

where $A = \mathbb{E}\left[\phi(\mathbf{s}_t)(\phi(\mathbf{s}_t) - \gamma\phi(\mathbf{s}'_t))^\top\right]$ and $b = \mathbb{E}\left[\mathbf{r}_t\phi(\mathbf{s}_t)\right]$. This gives rise to the *Least Squares Temporal Difference (LSTD)* algorithm [11, 12], which at each iteration t, provides estimates A_t of matrix A and $\mathbf{b}_t$ of vector b, and upon termination of the algorithm at time T, the approximation vector $\mathbf{z}_T$ is evaluated by $\mathbf{z}_T = (A_T)^{-1}\mathbf{b}_T$.

Least Squares Policy Evaluation (LSPE) [13] is a multi-stage algorithm where in the first stage, it obtains $\mathbf{u}_{t+1} = \arg\min_u \|\Phi\mathbf{z}_t - T^\pi\Phi u\|^2_\nu$ using the least squares method. In the subsequent stage, it minimizes the fix-point error using the recursion $\mathbf{z}_{t+1} = \mathbf{z}_t + \alpha_{t+1}\left(\mathbf{u}_{t+1} - \mathbf{z}_t\right)$.

Van Roy and Tsitsiklis [15] gave a different characterization for the limit point z^*_{TD} of TD(0) as the fixed point of the *projected Bellman operator* $\Pi^\nu T^\pi$,

$$\Phi z = \Pi^\nu T^\pi \Phi z. \tag{6}$$

This characterization yields a new error function, the *Mean Squared Projected Bellman Error (MSPBE)* defined as

$$\text{MSPBE}(z) \triangleq \|\Phi z - \Pi^\nu T^\pi \Phi z\|^2_\nu, \quad z \in \mathbb{R}^k. \tag{7}$$

In [9, 8], this objective function is maneuvered to derive novel $\Theta(k)$ algorithms like *GTD*, *TDC* and *GTD2*. GTD2 is a multi-timescale algorithm given by the following recursions:

$$\mathbf{z}_{t+1} = \mathbf{z}_t + \alpha_{t+1}\left(\phi(\mathbf{s}_t) - \gamma\phi(\mathbf{s}'_t)\right)\left(\phi(\mathbf{s}_t)^\top\mathbf{v}_t\right), \tag{8}$$

$$\mathbf{v}_{t+1} = \mathbf{v}_t + \beta_{t+1}(\delta_t(\mathbf{z}_t) - \phi(\mathbf{s}_t)^\top\mathbf{v}_t)\phi(\mathbf{s}_t). \tag{9}$$

The learning rates α_t and β_t satisfy $\sum_t \alpha_t = \infty$, $\sum_t \alpha_t^2 < \infty$ and $\beta_t = \eta\alpha_t$, where $\eta > 0$.

Another pertinent error function is the *Mean Square Bellman Residue (*MSBR*)* which is defined as

$$\text{MSBR}(z) \triangleq \mathbb{E}\left[\left(\mathbb{E}\left[\delta_t(z)|\mathbf{s}_t\right]\right)^2\right], z \in \mathbb{R}^k. \tag{10}$$

MSBR is a measure of how closely the prediction vector represents the solution to the Bellman equation.

Algorithm	Complexity	Error	Eligibility Trace
LSTD	$\Theta(k^3)$	MSPBE	Yes
TD	$\Theta(k)$	MSPBE	Yes
LSPE	$\Theta(k^3)$	MSPBE	Yes
GTD	$\Theta(k)$	MSPBE	-
GTD2	$\Theta(k)$	MSPBE	-
RG	$\Theta(k)$	MSBR	Yes

Table 1 Comparison of the state-of-the-art function approximation RL algorithms

Residual Gradient (RG) algorithm [10] minimizes the error function MSBR directly using stochastic gradient search. RG however requires double sampling, *i.e.*, generating two independent samples $\mathbf{s}_t'$ and $\mathbf{s}_t''$ of the next state when in the current state $\mathbf{s}_t$. The recursion is given by

$$\mathbf{z}_{t+1} = \mathbf{z}_t + \alpha_{t+1}\left(\mathbf{r}_t + \gamma \mathbf{z}_t^\top \phi_t' - \mathbf{z}_t^\top \phi_t\right)\left(\phi_t - \gamma \phi_t''\right), \tag{11}$$

where $\phi_t \triangleq \phi(\mathbf{s}_t)$, $\phi_t' \triangleq \phi(\mathbf{s}_{t+1}')$ and $\phi_t'' \triangleq \phi(\mathbf{s}_t'')$. Even though RG algorithm guarantees convergence, due to large variance, the convergence rate is small.

If the feature set, *i.e.*, the columns of the feature matrix Φ is linearly independent, then both the error functions MSBR and MSPBE are strongly convex. However, their respective minima are related depending on whether the feature set is perfect or not. A feature set is *perfect* if $V^\pi \in \{\Phi z | z \in \mathbb{R}^k\}$. If the feature set is perfect, then the respective minima of MSBR and MSPBE are the same. In the imperfect case, they differ. A relationship between MSBR and MSPBE can be easily established as follows:

$$\mathrm{MSBR}(z) = \mathrm{MSPBE}(z) + \|T^\pi \Phi z - \Pi^\nu T^\pi \Phi z\|^2, \quad z \in \mathbb{R}^k. \tag{12}$$

A vivid depiction of the relationship is shown in Figure 1.

Another relevant error objective is the *Mean Square Error (MSE)* which is the square of the ν-weighted distance from V^π and is defined as

$$\mathrm{MSE}(z) \triangleq \|V^\pi - \Phi z\|_\nu^2, \quad z \in \mathbb{R}^k. \tag{13}$$

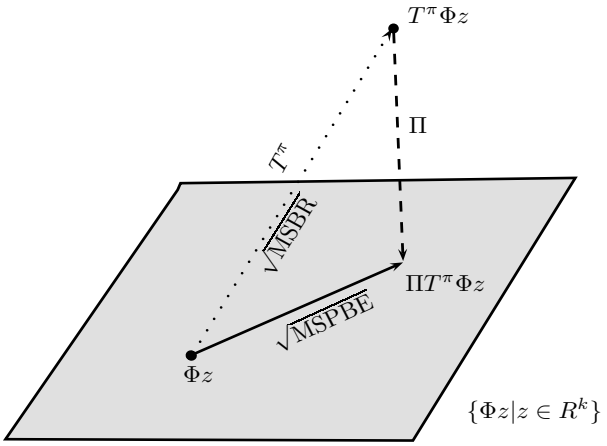

Figure 1 Diagram depicting the relationship between the error functions MSPBE and MSBR.

In [16] and [17] the relationship between MSE and MSBR is provided. It is found that, for a given ν with $\nu(s) > 0, \forall s \in \mathbb{S}$,

$$\sqrt{\text{MSE}(z)} \leq \frac{\sqrt{C(\nu)}}{1-\gamma} \sqrt{\text{MSBR}(z)}, \tag{14}$$

where $C(\nu) = \max_{s,s'} \frac{\text{P}(s,s')}{\nu(s)}$. Another bound which is of considerable importance is the bound on the MSE of the limit point z^*_{TD} of the TD(0) algorithm provided in [15]. It is found that

$$\sqrt{\text{MSE}(z^*_{\text{TD}})} \leq \frac{1}{\sqrt{1-\gamma^2}} \sqrt{\text{MSE}(z^\nu)}, \tag{15}$$

where $z^\nu \in \mathbb{R}^k$ satisfies $\Phi z^\nu = \Pi^\nu V^\pi$ and γ is the discount factor. Table 1 provides a list of important TD based algorithms along with the associated error objectives. The algorithm complexities are also shown in the table.

Put succinctly, when linear function approximation is applied in an RL setting, the main task can be cast as an optimization problem whose objective function is one of the aforementioned error functions. Typically, almost all the state-of-the-art algorithms employ gradient search technique to solve the minimization problem. In this paper, we apply a gradient-free technique called the *Cross Entropy (CE) method* instead to find the minimum. By '*gradient-free*', we mean the algorithm does not incorporate information on the gradient of the objective function, rather uses the function values themselves. Cross Entropy method is commonly subsumed within the general class of *Model based search methods* [18]. Other methods in this class

are *Model Reference Adaptive Search (MRAS)* [19], *Gradient-based Adaptive Stochastic Search for Simulation Optimization (GASSO)* [20], *Ant Colony Optimization (ACO)* [21] and *Estimation of Distribution Algorithms (EDAs)* [22]. Model based search methods have been applied to the control problem[1] in [23] and in basis adaptation[2] [24], but this is the first time such a procedure has been applied to the prediction problem. However, due to certain limitations in the original CE method, it cannot be directly applied to the RL setting. In this paper, we have proposed a method to workaround these limitations of the CE method, thereby making it a good choice for the RL setting. Note that any of the aforementioned error functions can be employed, but in this paper, we attempt to minimize MSPBE as it offers the best approximation with less bias to the projection $\Pi^{\nu} V^{\pi}$ for a given policy π, using a single sample trajectory.

Our Contributions The *Cross Entropy (CE) method* [25, 26] is a model based search algorithm to find the global maximum of a given real valued objective function. In this paper, we propose for the first time, an adaptation of this method to the problem of parameter tuning in order to find the best estimates of the value function V^{π} for a given policy π under the linear function approximation architecture. We propose a multi-timescale stochastic approximation algorithm which minimizes the MSPBE. The algorithm possesses the following attractive features:

1. No restriction on the feature set.
2. The computational complexity is quadratic in the number of features (this is a significant improvement compared to the cubic complexity of the least squares algorithms).
3. It is competitive with least squares and other state-of-the-art algorithms in terms of accuracy.
4. It is online with incremental updates.
5. It gives guaranteed convergence to the global minimum of the MSPBE.

A noteworthy observation is that since MSPBE is a strongly convex function [14], local and global minima overlap and the fact that CE method finds the global minima as opposed to local minima, unlike gradient search, is not really essential. Nonetheless, in the case of non-linear function approximators, the convexity property does not hold in general and so there may exist multiple local minima in the objective and the gradient search schemes would get stuck in local optima unlike CE based search. We have not explored the non-linear case in this paper. However, our approach can be viewed as a significant first step towards efficiently using model based search for policy evaluation in the RL setting.

2. Proposed Algorithm: SCE-MSPBEM

We present in this section our algorithm SCE-MSPBEM, acronym for *Stochastic Cross Entropy-Mean Squared Projected Bellman Error Minimization* that minimizes the Mean Squared Projected Bellman Error (MSPBE) by incorporating a multi-timescale stochastic approximation variant of the Cross Entropy (CE) method.

2.1. Summary of Notation:

Let $\mathbb{I}_{k\times k}$ and $0_{k\times k}$ be the identity matrix and the zero matrix with dimensions $k \times k$ respectively. Let $f_\theta(\cdot)$ be the probability density function (*pdf*) parametrized by θ and $\mathbb{P}_\theta$ be its induced probability measure. Let $\mathbb{E}_\theta[\cdot]$ be the expectation *w.r.t.* the probability distribution $f_\theta(\cdot)$. We define the $(1-\rho)$-quantile of a real-valued function $\mathcal{H}(\cdot)$ *w.r.t.* the probability distribution $f_\theta(\cdot)$ as follows:

$$\gamma_\rho^{\mathcal{H}}(\theta) \triangleq \sup\{l \in \mathbb{R} : \mathbb{P}_\theta(\mathcal{H}(\mathbf{x}) \geq l) \geq \rho\}, \tag{16}$$

Also $\lceil a \rceil$ denotes the smallest integer greater than a. For $A \subset \mathbb{R}^m$, let I_A represent the indicator function, *i.e.*, $I_A(x)$ is 1 when $x \in A$ and is 0 otherwise. We denote by $\mathbb{Z}_+$ the set of non-negative integers. Also we denote by $\mathbb{R}_+$ the set of non-negative real numbers. Thus, 0 is an element of both $\mathbb{Z}_+$ and $\mathbb{R}_+$. In this section, $\mathbf{x}$ represents a random variable and x a deterministic variable.

2.2. Background: The CE Method

To better understand our algorithm, we briefly explicate the original CE method first.

2.2.1. Objective of CE The *Cross Entropy (CE) method* [25, 27, 26] solves problems of the following form:

$$\text{Find} \quad x^* \in \arg\max_{x \in \mathcal{X} \subset \mathbb{R}^m} \mathcal{H}(x),$$

where $\mathcal{H}(\cdot)$ is a multi-modal real-valued function and $\mathcal{X}$ is called the solution space.

The goal of the CE method is to find an optimal "*model*" or probability distribution over the solution space $\mathcal{X}$ which concentrates on the global maxima of $\mathcal{H}(\cdot)$. The CE method adopts an iterative procedure where at each iteration t, a search is conducted on a space of parametrized probability distributions $\{f_\theta | \theta \in \Theta\}$ on

$\mathcal{X}$ over $\mathcal{X}$, where Θ is the parameter space, to find a distribution parameter θ_t which reduces the *Kullback-Leibler (KL)* distance from the optimal model. The most commonly used class here is the *exponential family of distributions*.

Exponential Family of Distributions: These are denoted as $\mathcal{C} \triangleq \{f_\theta(x) = h(x)e^{\theta^\top \Gamma(x) - K(\theta)} \mid \theta \in \Theta \subset \mathbb{R}^d\}$, where $h : \mathbb{R}^m \longrightarrow \mathbb{R}$, $\Gamma : \mathbb{R}^m \longrightarrow \mathbb{R}^d$ and $K : \mathbb{R}^d \longrightarrow \mathbb{R}$. By rearranging the parameters, we can show that the Gaussian distribution with mean vector μ and the covariance matrix Σ belongs to $\mathcal{C}$. In this case,

$$f_\theta(x) = \frac{1}{\sqrt{(2\pi)^m |\Sigma|}} e^{-(x-\mu)^\top \Sigma^{-1}(x-\mu)/2}, \tag{17}$$

and so one may let $h(x) = \frac{1}{\sqrt{(2\pi)^m}}$, $\Gamma(x) = (x, xx^\top)^\top$ and $\theta = (\Sigma^{-1}\mu,\ -\frac{1}{2}\Sigma^{-1})^\top$.

⊛ **Assumption (A1):** The parameter space Θ is compact.

2.2.2. CE Method (Ideal Version) The CE method aims to find a sequence of model parameters $\{\theta_t\}_{t\in\mathbb{Z}^+}$, where $\theta_t \in \Theta$ and an increasing sequence of thresholds $\{\gamma_t\}_{t\in\mathbb{Z}^+}$ where $\gamma_t \in \mathbb{R}$, with the property that the support of the model identified using θ_t, *i.e.*, $\overline{\{x|f_{\theta_t}(x) \neq 0\}}$ is contained in the region $\{x|\mathcal{H}(x) \geq \gamma_t\}$. By assigning greater weight to higher values of $\mathcal{H}$ at each iteration, the expected behaviour of the probability distribution sequence should improve. The most common choice for γ_{t+1} is $\gamma_\rho^{\mathcal{H}}(\theta_t)$, the $(1-\rho)$-quantile of $\mathcal{H}(\mathbf{x})$ *w.r.t.* the probability distribution $f_{\theta_t}(\cdot)$, where $\rho \in (0,1)$ is set *a priori* for the algorithm. We take Gaussian distribution as the preferred choice for $f_\theta(\cdot)$ in this paper. In this case, the model parameter is $\theta = (\mu, \Sigma)^\top$ where $\mu \in \mathbb{R}^m$ is the mean vector and $\Sigma \in \mathbb{R}^{m\times m}$ is the covariance matrix.

The CE algorithm is an iterative procedure which starts with an initial value $\theta_0 = (\mu_0, \Sigma_0)^\top$ of the mean vector and the covariance matrix tuple and at each iteration t, a new parameter $\theta_{t+1} = (\mu_{t+1}, \Sigma_{t+1})^\top$ is derived from the previous value θ_t as follows:

$$\theta_{t+1} = \arg\max_{\theta\in\Theta} \mathbb{E}_{\theta_t}\left\{S(\mathcal{H}(\mathbf{x}))I_{\{\mathcal{H}(\mathbf{x})\geq\gamma_{t+1}\}} \log f_\theta(\mathbf{x})\right\}, \tag{18}$$

where S is positive and strictly monontone.

If the gradient *w.r.t.* θ of the objective function in (18) is equated to 0 and using (17) for $f_\theta(\cdot)$, we obtain

$$\mu_{t+1} = \frac{\mathbb{E}_{\theta_t}\left[\mathbf{g_1}(\mathcal{H}(\mathbf{x}), \mathbf{x}, \gamma_{t+1})\right]}{\mathbb{E}_{\theta_t}\left[\mathbf{g_0}(\mathcal{H}(\mathbf{x}), \gamma_{t+1})\right]} \triangleq \Upsilon_1(\mathcal{H}(\cdot), \theta_t, \gamma_{t+1}), \tag{19}$$

$$\Sigma_{t+1} = \frac{\mathbb{E}_{\theta_t}\left[\mathbf{g_2}(\mathcal{H}(\mathbf{x}), \mathbf{x}, \gamma_{t+1}, \mu_{t+1})\right]}{\mathbb{E}_{\theta_t}\left[\mathbf{g_0}(\mathcal{H}(\mathbf{x}), \gamma_{t+1})\right]} \triangleq \Upsilon_2(\mathcal{H}(\cdot), \theta_t, \gamma_{t+1}, \mu_{t+1}). \tag{20}$$

where $\mathbf{g_0}(\mathcal{H}(x), \gamma) \triangleq S(\mathcal{H}(x)) I_{\{\mathcal{H}(x) \geq \gamma\}}$, (21)

$$\mathbf{g_1}(\mathcal{H}(x), x, \gamma) \triangleq S(\mathcal{H}(x)) I_{\{\mathcal{H}(x) \geq \gamma\}} x, \tag{22}$$

$$\mathbf{g_2}(\mathcal{H}(x), x, \gamma, \mu) \triangleq S(\mathcal{H}(x)) I_{\{\mathcal{H}(x) \geq \gamma\}} (x-\mu)(x-\mu)^\top. \tag{23}$$

REMARK 1. The function $S(\cdot)$ in (18) is positive and strictly monotone and is used to account for the cases when the objective function $\mathcal{H}(x)$ takes negative values for some x. One common choice is $S(x) = \exp(rx)$ where $r \in \mathbb{R}$ is chosen appropriately.

2.2.3. CE Method (Monte-Carlo Version) It is hard in general to evaluate $\mathbb{E}_{\theta_t}[\cdot]$ and γ_t, so the stochastic counterparts of the equations (19) and (20) are used instead in the CE algorithm. This gives rise to the Monte-Carlo version of the CE method. In this stochastic version, the algorithm generates a sequence $\{\bar{\theta}_t = (\bar{\mu}_t, \bar{\Sigma}_t)^\top\}_{t \in \mathbb{Z}_+}$, where at each iteration t, N_t samples $\Lambda_t = \{\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_{N_t}\}$ are picked using the distribution $f_{\bar{\theta}_t}$ and the estimate of γ_{t+1} is obtained as follows: $\bar{\gamma}_{t+1} = \mathcal{H}_{(\lceil (1-\rho) N_t \rceil)}$ where $\mathcal{H}_{(i)}$ is the ith-order statistic of $\{\mathcal{H}(\mathbf{x}_i)\}_{i=1}^{N_t}$. The estimate $\bar{\theta}_{t+1} = (\bar{\mu}_{t+1}, \bar{\Sigma}_{t+1})^\top$ of the model parameters $\theta_{t+1} = (\mu_{t+1}, \Sigma_{t+1})^\top$ is obtained as

$$\bar{\mu}_{t+1} = \frac{\frac{1}{N_t}\sum_{i=1}^{N_t} \mathbf{g_1}(\mathcal{H}(\mathbf{x}_i), \mathbf{x}_i, \bar{\gamma}_{t+1})}{\frac{1}{N_t}\sum_{i=1}^{N_t} \mathbf{g_0}(\mathcal{H}(\mathbf{x}_i), \bar{\gamma}_{t+1})} \triangleq \bar{\Upsilon}_1(\mathcal{H}(\cdot), \bar{\theta}_t, \bar{\gamma}_{t+1}), \tag{24}$$

$$\bar{\Sigma}_{t+1} = \frac{\frac{1}{N_t}\sum_{i=1}^{N_t} \mathbf{g_2}(\mathcal{H}(\mathbf{x}_i), \mathbf{x}_i, \bar{\gamma}_{t+1}, \bar{\mu}_{t+1})}{\frac{1}{N_t}\sum_{i=1}^{N_t} \mathbf{g_0}(\mathcal{H}(\mathbf{x}_i), \bar{\gamma}_{t+1})} \triangleq \bar{\Upsilon}_2(\mathcal{H}(\cdot), \bar{\theta}_t, \bar{\gamma}_{t+1}, \bar{\mu}_{t+1}). \tag{25}$$

An *observation allocation rule* $\{N_t \in \mathbb{Z}_+\}_{t\in\mathbb{Z}_+}$ is used to determine the sample size. The Monte-Carlo version of the CE method is described in Algorithm 1.

Algorithm 1: The Monte-Carlo CE Algorithm

Step 0: Choose an initial *p.d.f.* $f_{\bar{\theta}_0}(\cdot)$ on $\mathcal{X}$ where $\bar{\theta}_0 = (\bar{\mu}_0, \bar{\Sigma}_0)^\top$ and fix an $\epsilon > 0$;

Step 1: [Sampling Candidate Solutions] Randomly sample N_t independent and identically distributed solutions $\Lambda_t = \{\mathbf{x}_1, \ldots, \mathbf{x}_{N_t}\}$ using $f_{\bar{\theta}_t}(\cdot)$.

Step 2: [Threshold Evaluation] Calculate the sample $(1-\rho)$-quantile $\bar{\gamma}_{t+1} = \mathcal{H}_{(\lceil (1-\rho)N_t \rceil)}$, where $\mathcal{H}_{(i)}$ is the ith-order statistic of the sequence $\{\mathcal{H}(\mathbf{x}_i)\}_{i=1}^{N_t}$;

Step 3: [Threshold Comparison]

if $\bar{\gamma}_{t+1} \geq \bar{\gamma}_t^* + \epsilon$ **then**
$\bar{\gamma}_{t+1}^* = \bar{\gamma}_{t+1}$,

else
$\bar{\gamma}_{t+1}^* = \bar{\gamma}_t^*$.

Step 3: [Model Parameter Update]

$\bar{\theta}_{t+1} = (\bar{\mu}_{t+1}, \bar{\Sigma}_{t+1})^\top = \left(\bar{\Upsilon}_1(\mathcal{H}(\cdot), \bar{\theta}_t, \bar{\gamma}_{t+1}^*), \bar{\Upsilon}_2(\mathcal{H}(\cdot), \bar{\theta}_t, \bar{\gamma}_{t+1}^*, \bar{\mu}_{t+1})\right)^\top$.

Step 4: If the stopping rule is satisfied, then return $\bar{\theta}_{t+1}$, else set $t := t+1$ and go to Step 1.

REMARK 2. The CE method is also applied in stochastic settings for which the objective function is given by $\mathcal{H}(x) = \mathbb{E}_{\mathbf{y}}[G(x, \mathbf{y})]$, where $\mathbf{y} \in \mathcal{Y}$ and $\mathbb{E}_{\mathbf{y}}[\cdot]$ is the expectation w.r.t. a probability distribution on $\mathcal{Y}$. Since the objective function is expressed in terms of expectation, it might be hard in some scenarios to obtain the true values of the objective function. In such cases, estimates of the objective function are used instead. The CE method is shown to have global convergence properties in such cases too.

2.2.4. Limitations of the CE Method A significant limitation of the CE method is its dependence on the sample size N_t used in Step 1 of Algorithm 1. One does not know *a priori the best value for the sample size* N_t. Higher values of N_t while resulting in higher accuracy also require more computational resources. One often needs to apply brute force in order to obtain a good choice of N_t. Also as m, the dimension of the solution space, takes large values, more samples are required for better accuracy, making N_t *large* as well. This makes *finding the ith-order statistic $\mathcal{H}_{(i)}$ in Step 2 harder*. Note that the order statistic $\mathcal{H}_{(i)}$ is obtained by sorting the list $\{\mathcal{H}(\mathbf{x}_1), \mathcal{H}(\mathbf{x}_2), \ldots \mathcal{H}(\mathbf{x}_{N_t})\}$. The computational effort required in that case is

$O(N_t \log N_t)$ which in most cases is inadmissible. The other major bottleneck is the *space required to store the samples* Λ_t. In situations when m and N_t are large, the storage requirement is a major concern.

The CE method is also *offline in nature.* This means that the function values $\{\mathcal{H}(\mathbf{x}_1), \ldots, \mathcal{H}(\mathbf{x}_{N_t})\}$ of the sample set $\Lambda_t = \{\mathbf{x}_1, \ldots, \mathbf{x}_{N_t}\}$ should be available before the model parameters can be updated in Step 4 of Algorithm 1. So when applied in the prediction problem of approximating the value function V^π for a given policy π using the linear architecture defined in (3) by minimizing the error function MSPBE$(\cdot)$, we require the estimates of $\{\text{MSPBE}(\mathbf{x}_1), \ldots, \text{MSPBE}(\mathbf{x}_{N_t})\}$. This means that a sufficiently long traversal along the given sample trajectory has to be conducted to obtain the estimates before initiating the CE method. This does not make the CE method amenable to online implementations in RL, where the value function estimations are performed in real-time after each observation.

In this paper, we resolve all these shortcomings of the CE method by remodelling the same in the stochastic approximation framework and thus *replacing the sample averaging operation in equations (24) and (25) with a bootstrapping approach* where we continuously improve the estimates based on past observations. Each successive estimate is obtained as a function of the previous estimate and a noise term. We replace the $(1-\rho)$-quantile estimation using the order statistic method in Step 2 of Algorithm 1 with a stochastic recursion which serves the same purpose, but more efficiently. The model parameter update in step 3 is also replaced with a stochastic recursion. We also bring in additional modifications to the CE method to adapt to a Markov Reward Process (MRP) framework and thus obtain an *online version* of CE where the computational requirements are quadratic in the size of the feature set for each observation. To fit the online nature, we have developed an expression for the objective function MSPBE, where we are able to separate its deterministic and non-deterministic components. This separation is critical since the original expression of MSPBE is convoluted with the solution vector and the expectation terms and hence is unrealizable. The separation further helps to develop a stochastic recursion for estimating MSPBE. Finally, in this paper, we provide a proof of convergence of our algorithm using an ODE based analysis.

2.3. Proposed Algorithm (SCE-MSPBEM)

Notation: In this section, $\mathbf{z}$ represents a random variable and z a deterministic variable.

SCE-MSPBEM is an *algorithm to approximate the value function V^π (for a given policy π) with linear*

function approximation, where the optimization is performed using a multi-timescale stochastic approximation variant of the CE algorithm. Since the CE method is a maximization algorithm, the objective function in the optimization problem here is the negative of MSPBE. Thus,

$$z^* = \arg\min_{z \in \mathcal{Z} \subset \mathbb{R}^k} \text{MSPBE}(z) = \arg\max_{z \in \mathcal{Z} \subset \mathbb{R}^k} \mathcal{J}(z), \quad \text{where } \mathcal{J} = -\text{MSPBE}. \tag{26}$$

Here $\mathcal{Z}$ is the solution space, *i.e.*, the space of parameter values of the function approximator. We also define $\mathcal{J}^* \triangleq \mathcal{J}(z^*)$.

REMARK 3. Since $\exists z \in \mathcal{Z}$ such that $\Phi z = \Pi^\nu T^\pi \Phi z$, the value of $\mathcal{J}^*$ is 0.

⊛ **Assumption (A2):** The solution space $\mathcal{Z}$ is compact, *i.e.*, it is closed and bounded.

In [9] a compact expression for MSPBE is given as follows:

$$\text{MSPBE}(z) = (\Phi^\top D^\nu (T_\pi V_z - V_z))^\top (\Phi^\top D^\nu \Phi)^{-1} (\Phi^\top D^\nu (T_\pi V_z - V_z)). \tag{27}$$

Using the fact that $V_z = \Phi z$, the expression $\Phi^\top D^\nu (T_\pi V_z - V_z)$ can be rewritten as

$$\begin{aligned} \Phi^\top D^\nu (T_\pi V_z - V_z) &= \mathbb{E}\left[\mathbb{E}\left[\phi_t(\mathbf{r}_t + \gamma z^\top \phi_t' - z^\top \phi_t)|\mathbf{s}_t\right]\right] \\ &= \mathbb{E}\left[\mathbb{E}\left[\phi_t \mathbf{r}_t|\mathbf{s}_t\right]\right] + \mathbb{E}\left[\mathbb{E}\left[\phi_t(\gamma\phi_t' - \phi_t)^\top|\mathbf{s}_t\right]\right] z, \quad \text{where } \phi_t \triangleq \phi(\mathbf{s}_t) \text{ and } \phi_t' \triangleq \phi(\mathbf{s}_t'). \end{aligned} \tag{28}$$

$$\text{Also,} \qquad \Phi^\top D^\nu \Phi = \mathbb{E}\left[\phi_t \phi_t^\top\right]. \tag{29}$$

Putting all together we get,

$$\begin{aligned} \text{MSPBE}(z) = \left(\mathbb{E}\left[\mathbb{E}\left[\phi_t \mathbf{r}_t|\mathbf{s}_t\right]\right] + \mathbb{E}\left[\mathbb{E}\left[\phi_t(\gamma\phi_t' - \phi_t)^\top|\mathbf{s}_t\right]\right] z\right)^\top (\mathbb{E}\left[\phi_t \phi_t^\top\right])^{-1} \\ \left(\mathbb{E}\left[\mathbb{E}\left[\phi_t \mathbf{r}_t|\mathbf{s}_t\right]\right] + \mathbb{E}\left[\mathbb{E}\left[\phi_t(\gamma\phi_t' - \phi_t)^\top|\mathbf{s}_t\right]\right] z\right). \end{aligned} \tag{30}$$

$$= \left(\omega_*^{(0)} + \omega_*^{(1)} z\right)^\top \omega_*^{(2)} \left(\omega_*^{(0)} + \omega_*^{(1)} z\right), \tag{31}$$

where $\omega_*^{(0)} \triangleq \mathbb{E}\left[\mathbb{E}\left[\phi_t \mathbf{r}_t|\mathbf{s}_t\right]\right]$, $\omega_*^{(1)} \triangleq \mathbb{E}\left[\mathbb{E}\left[\phi_t(\gamma\phi_t' - \phi_t)^\top|\mathbf{s}_t\right]\right]$ and $\omega_*^{(2)} \triangleq (\mathbb{E}\left[\phi_t \phi_t^\top\right])^{-1}$.

This is a quadratic function on z. Note that, in the above expression, the parameter vector z and the stochastic component involving $\mathbb{E}[\cdot]$ are decoupled. Hence the stochastic component can be estimated independent of the parameter vector z.

The goal of this paper is to adapt CE method into a MRP setting in an online fashion, where we solve the prediction problem which is a continuous stochastic optimization problem. The important tool we employ here to achieve this is the stochastic approximation framework. Here we take a slight digression to explain stochastic approximation algorithms.

Stochastic approximation algorithms [28, 29, 30] are a natural way of utilizing prior information. It does so by discounted averaging of the prior information and are usually expressed as recursive equations of the following form:

$$\mathbf{z}_{j+1} = \mathbf{z}_j + \alpha_{j+1}\Delta\mathbf{z}_{j+1}, \tag{32}$$

where $\Delta\mathbf{z}_{j+1} = q(\mathbf{z}_j) + b_j + \mathbb{M}_{j+1}$ is the *increment term*, $q(\cdot)$ is a Lipschitz continuous function, b_j is the *bias term* with $b_j \to 0$ and $\{\mathbb{M}_j\}$ is a *martingale difference noise sequence*, *i.e.*, $\mathbb{M}_j$ is $\mathcal{F}_j$-measurable and integrable and $\mathbb{E}[\mathbb{M}_{j+1}|\mathcal{F}_j] = 0, \forall j$. Here $\{\mathcal{F}_j\}_{j\in\mathbb{Z}_+}$ is a filtration, where the σ-field $\mathcal{F}_j = \sigma(\mathbf{z}_i, \mathbb{M}_i, 1 \leq i \leq j, \mathbf{z}_0)$. The learning rate α_j satisfies $\Sigma_j\alpha_j = \infty$, $\Sigma_j\alpha_j^2 < \infty$.

We have the following well known result from [28] regarding the limiting behaviour of the stochastic recursion (32):

THEOREM 1. *Assume* $\sup_j \|z_j\| < \infty$, $\mathbb{E}[\|\mathbb{M}_{j+1}\|^2|\mathcal{F}_j] \leq K(1 + \|\mathbf{z}_j\|^2), \forall j$ *and* $q(\cdot)$ *is Lipschitz continuous. Then the iterates* $\mathbf{z}_j$ *converge almost surely to the compact connected internally chain transitive invariant set of the ODE:*

$$\dot{z}(t) = q(z(t)), t \in \mathbb{R}_+. \tag{33}$$

Put succinctly, the above theorem establishes an equivalence between the asymptotic behaviour of the iterates $\{\mathbf{z}_j\}$ and the deterministic ODE (33). In most practical cases, the ODE have a unique globally asymptotically stable equilibrium at an arbitrary point z^*. It will then follow from the theorem that $\mathbf{z}_j \to z^*$ *a.s.* However, in some cases, the ODE can have multiple isolated stable equilibria. In such cases, the convergence of $\{\mathbf{z}_j\}$ to one of these equilibria is guaranteed, however the limit point would depend on the noise and the initial value.

A relevant extension of the stochastic approximation algorithms is the multi-timescale variant. Here there will be multiple stochastic recursions of the kind (32), each with possibly different learning rates. The

learning rates defines the timescale of the particular recursion. So different learning rates imply different timescales. If the increment terms are well-behaved and the learning rates properly related (defined in Chapter 6 of [28]), then the chain of recursions exhibit a well-defined asymptotic behaviour. See Chapter 6 [28] for more details.

Now digression aside, note that the important stochastic variables of the ideal CE method are $\mathcal{H}, \gamma_k$, μ_k, Σ_k and θ_k. Here, the objective function $\mathcal{H} = \mathcal{J}$. In our approach, we track these variables independently using stochastic recursions of the kind (32). Thus we model our algorithm as a multi-timescale stochastic approximation algorithm which tracks the ideal CE method. Note that the stochastic recursion is uniquely identified by their increment term, their initial value and the learning rate. We consider here these recursions in great detail.

1. Tracking the Objective Function $\mathcal{J}(\cdot)$: Recall that the goal of the paper is to develop an online and incremental prediction algorithm. This implies that algorithm has to learn from a given sample trajectory using an incremental approach with a single traversal of the trajectory. The algorithm SCE-MSPBEM operates online on a single sample trajectory $\{(\mathbf{s}_t, \mathbf{r}_t, \mathbf{s}'_t)\}_{t=0}^{\infty}$, where $\mathbf{s}_t \sim \nu(\cdot)$, $\mathbf{s}'_t \sim \mathrm{P}^{\pi}(\mathbf{s}_t, \cdot)$ and $\mathbf{r}_t = \mathrm{R}(\mathbf{s}_t, \pi(\mathbf{s}_t), \mathbf{s}'_t)$.

⊛ **Assumption (A3):** For the given trajectory $\{(\mathbf{s}_t, \mathbf{r}_t, \mathbf{s}'_t)\}_{t=0}^{\infty}$, let ϕ_t, ϕ'_t, and $\mathbf{r}_t$ have uniformly bounded second moments. Also, $\mathbb{E}\left[\phi_t \phi_t^{\top}\right]$ is non-singular.

In the expression (31) for the objective function $\mathcal{J}(\cdot)$, we have isolated the stochastic and deterministic part. The stochastic part can be identified by the tuple $\omega_* \triangleq (\omega_*^{(0)}, \omega_*^{(1)}, \omega_*^{(2)})^{\top}$. So if we can find ways to track ω_*, then it implies we could track $\mathcal{J}(\cdot)$. This is the line of thought we follow here. In our algorithm, we track ω_* using the time dependent variable $\omega_t \triangleq (\omega_t^{(0)}, \omega_t^{(1)}, \omega_t^{(2)})^{\top}$, where $\omega_t^{(0)} \in \mathbb{R}^k$, $\omega_t^{(1)} \in \mathbb{R}^{k \times k}$ and $\omega_t^{(2)} \in \mathbb{R}^{k \times k}$. Here $\omega_t^{(i)}$ independently tracks $\omega_*^{(i)}$, $1 \leq i \leq 3$. Note that tracking implies $\lim_{t \to \infty} \omega_t^{(i)} = \omega_*^{(i)}$, $1 \leq i \leq 3$. The increment term $\Delta\omega_t \triangleq (\omega_t^{(0)}, \omega_t^{(1)}, \omega_t^{(2)})^{\top}$ used for this recursion is defined as follows:

$$\left.\begin{aligned} \triangle\omega_t^{(0)} &= \mathbf{r}_t \phi_t - \omega_t^{(0)}, \\ \triangle\omega_t^{(1)} &= \phi_t(\gamma\phi'_t - \phi_t)^{\top} - \omega_t^{(1)}, \\ \triangle\omega_t^{(2)} &= \mathbb{I}_{k \times k} - \phi_t \phi_t^{\top} \omega_t^{(2)}, \end{aligned}\right\} \tag{34}$$

where $\phi_t \triangleq \phi(\mathbf{s}_t)$ and $\phi'_t \triangleq \phi(\mathbf{s}'_t)$. Now we define a new function $\bar{\mathcal{J}}(\omega_t, z) \triangleq -\left(\omega_t^{(0)} + \omega_t^{(1)} z\right)^\top \omega_t^{(2)} \left(\omega_t^{(0)} + \omega_t^{(1)} z\right)$. Note that this is the same expression as (31) except for ω_t replacing ω_*. Since ω_t tracks ω_*, it is easily verifiable that $\bar{\mathcal{J}}(\omega_t, z)$ indeed tracks $\mathcal{J}(z)$ for a given $z \in \mathcal{Z}$.

The stochastic recursions which track ω_* and the objective function $\mathcal{J}(\cdot)$ are defined in (41) and (42) respectively. A rigorous analysis of the above stochastic recursion is provided in lemma 2. There we also find that the initial value ω_0 is irrelevant.

2. Tracking $\gamma_\rho(\mathcal{J}, \theta)$: Here we are faced two difficult situations: (i) the true objective function $\mathcal{J}$ is unavailable and (ii) we have to find a stochastic recursion which tracks $\gamma_\rho(\mathcal{J}, \theta)$ for a given distribution parameter θ. To solve (i) we use whatever is available, *i.e.* $\bar{\mathcal{J}}(\omega_t, \cdot)$ which is the best available estimate of the true function $\mathcal{J}(\cdot)$ at time t. In other words, we bootstrap. Now to address the second part we make use of the following lemma from [31]. The lemma provides a characterization of the $(1-\rho)$-quantile of a given real-valued function H *w.r.t.* to a given probability distribution function f_θ.

LEMMA 1. *The $(1-\rho)$-quantile of a bounded real valued function $H(\cdot)$ $\left(\text{with } H(x) \in [H_l, H_u]\right)$ w.r.t the probability density function $f_\theta(\cdot)$ is reformulated as an optimization problem*

$$\gamma_\rho(H, \theta) = \underset{y \in [H_l, H_u]}{\arg\min}\, \mathbb{E}_\theta \left[\psi(H(\mathbf{x}), y)\right], \tag{35}$$

where $\mathbf{x} \sim f_\theta(\cdot)$, $\psi(H(x), y) = (1-\rho)(H(x) - y) I_{\{H(x) \geq y\}} + \rho(y - H(x)) I_{\{H(x) \leq y\}}$ and $\mathbb{E}_\theta[\cdot]$ is the expectation w.r.t. the p.d.f. $f_\theta(\cdot)$.

In this paper, we employ the time-dependent variable γ_t to track $\gamma_\rho(\mathcal{J}, \cdot)$. The increment term in the recursion is the subdifferential $\nabla_y \psi$. This is because ψ is non-differentiable as it follows from its definition from the above lemma. However subdifferential exists for ψ. Hence we utilize it to solve the optimization problem (35) defined in lemma 1. Here, we define an increment function (contrary to an increment term) and is defined as follows:

$$\Delta\gamma_{t+1}(z) = -(1-\rho) I_{\{\bar{\mathcal{J}}(\omega_t, z) \geq \gamma_t\}} + \rho I_{\{\bar{\mathcal{J}}(\omega_t, z) \leq \gamma_t\}} \tag{36}$$

The stochastic recursion which tracks $\gamma_\rho(\mathcal{J},\cdot)$ is given in (43). A deeper analysis of the recursion (43) is provided in lemma 3. In the analysis we also find that the initial value γ_0 is irrelevant.

3. Tracking Υ_1 and Υ_2: In the ideal CE method, for a given θ_t, note that $\Upsilon_1(\theta_t,\dots)$ and $\Upsilon_2(\theta_t,\dots)$ form the subsequent model parameter θ_{t+1}. In our algorithm, we completely avoid the sample averaging technique employed in the Monte-Carlo version. Instead, we follow the stochastic approximation recursion to track the above quantities. Two time-dependent variables $\xi_t^{(1)}$ and $\xi_t^{(2)}$ are employed to track Υ_1 and Υ_2 respectively. The increment functions used by their respective recursions are defined as follows:

$$\Delta\xi_{t+1}^{(0)}(z) = \quad \mathbf{g_1}(\bar{\mathcal{J}}(\omega_t,z),z,\gamma_t) - \xi_t^{(0)}\mathbf{g_0}(\bar{\mathcal{J}}(\omega_t,z),\gamma_t), \tag{37}$$

$$\Delta\xi_{t+1}^{(1)}(z) = \mathbf{g_2}(\bar{\mathcal{J}}(\omega_t,z),z,\gamma_t,\xi_t^{(0)}) - \xi_t^{(1)}\mathbf{g_0}(\bar{\mathcal{J}}(\omega_t,z),\gamma_t). \tag{38}$$

The recursive equations which track Υ_1 and Υ_2 are defined in (44) and (45) respectively. The analysis of these recursions is provided in lemma 4. In this case also, the initial values are irrelevant.

4. Model Parameter Update: In the ideal version of CE, note that given θ_t, we have $\theta_{t+1} = (\Upsilon_1(\mathcal{H}(\cdot),\theta_t,\dots),\Upsilon_2(\mathcal{H}(\cdot),\theta_t,\dots))^\top$. This is a discrete change from θ_t to θ_{t+1}. But in our algorithm, we adopt a smooth update of the model parameters. The recursion is defined in equation (48). We prove in theorem 2 that the above approach indeed provide an optimal solution to the optimization problem defined in (26).

5. Learning Rates and Timescales: The algorithm uses two learning rates α_t and β_t which are deterministic, positive, nonincreasing and satisfy the following conditions:

$$\sum_{t=1}^{\infty}\alpha_t = \sum_{t=1}^{\infty}\beta_t = \infty, \qquad \sum_{t=1}^{\infty}\left(\alpha_t^2+\beta_t^2\right) < \infty, \qquad \lim_{t\to\infty}\frac{\alpha_t}{\beta_t} = 0. \tag{39}$$

In a multi-timescale stochastic approximation setting, it is important to understand the difference between timescale and learning rate. The timescale of a stochastic recursion is defined by its learning rate (also referred as step-size). Note that from the conditions imposed on the learning rates α_t and β_t in (39), we have $\frac{\alpha_t}{\beta_t}\to 0$. So α_t decays to 0 faster than β_t. Hence the timescale obtained from $\beta_t, t\geq 0$ is considered faster as compared to the other. So in a multi-timescale stochastic recursion scenario, the evolution of the recursion

controlled by the faster step-sizes (converges faster to 0) is slower compared to the recursions controlled by the slower step-sizes. This is because the increments are weighted by their learning rates, *i.e.*, the learning rates control the quantity of change that occurs to the variables when the update is executed. So the faster timescale recursions converge faster compared to its slower counterparts. Infact, when observed from a faster timescale recursion, one can consider the slower timescale recursion to be almost stationary. This attribute of the multi-timescale recursions are very important in the analysis of the algorithm. In the analysis, when studying the asymptotic behaviour of a particular stochastic recursion, we can consider the variables of other recursions which are on slower timescales to be constant. In our algorithm, the recursion of ω_t and θ_t proceed along the slowest timescale and so updates of ω_t appear to be quasi-static when viewed from the timescale on which the recursions governed by β_t proceed. The recursions of $\gamma_t, \xi_t^{(0)}$ and $\xi_t^{(1)}$ proceed along the faster timescale and hence have a faster convergence rate. The stable behaviour of the algorithm is attributed to the timescale differences obeyed by the various recursions.

6. Sample Requirement: The streamline nature inherent in the stochastic approximation algorithms demands only a single sample per iteration. Infact, we use two samples $\mathbf{z}_{t+1}$ (generated in (40)) and $\mathbf{z}_{t+1}^p$ (generated in (46) whose discussion is deferred for the time being).This is a remarkable improvement, apart from the fact that the algorithm is now online and incremental in the sense that whenever a new state transition $(\mathbf{s}_t, \mathbf{r}_t, \mathbf{s}_t')$ is revealed, the algorithm learns from it by evolving the variables involved and directing the model parameter θ_t towards the degenerate distribution concentrated on the optimum point z^*.

7. Mixture Distribution: In the algorithm, we use a mixture distribution $\widehat{f}_{\theta_t}$ to generate the sample $\mathbf{x}_{t+1}$, where $\widehat{f}_{\theta_t} = (1-\lambda) f_{\theta_t} + \lambda f_{\theta_0}$ with $\lambda \in (0,1)$ the mixing weight. The initial distribution parameter θ_0 is chosen *s.t.* the density function f_{θ_0} is strictly positive on every point in the solution space $\mathcal{X}$, *i.e.*, $f_{\theta_0}(x) > 0, \forall x \in \mathcal{X}$. The mixture approach facilitates exploration of the solution space and prevents the iterates from getting stranded in suboptimal solutions.

The SCE-MSPBEM algorithm is formally presented in Algorithm 2.

Algorithm 2: SCE-MSPBEM

Data: $\alpha_t, \beta_t, c_t \in (0,1)$, $c_t \to 0$, $\epsilon_1, \lambda, \rho \in (0,1)$, $S(\cdot): \mathbb{R} \to \mathbb{R}^+$;

Initialization: $\gamma_0 = 0$, $\gamma_0^p = -\infty$, $\theta_0 = (\mu_0, \Sigma_0)^\top$, $T_0 = 0$, $\xi_t^{(0)} = 0_{k\times 1}$, $\xi_t^{(1)} = 0_{k\times k}$,

$$\omega_0^{(0)} = 0_{k\times 1}, \omega_0^{(1)} = 0_{k\times k}, \omega_0^{(2)} = 0_{k\times k}, \theta^p = NULL;$$

foreach $(\mathbf{s}_t, \mathbf{r}_t, \mathbf{s}'_t)$ *of the trajectory* **do**

$$\left.\begin{aligned} \widehat{f}_{\theta_t} &= (1-\lambda) f_{\theta_t} + \lambda f_{\theta_0}; \\ \mathbf{z}_{t+1} &\sim \widehat{f}_{\theta_t}(\cdot); \end{aligned}\right\} \tag{40}$$

• [Objective Function Evaluation]

$$\omega_{t+1} = \omega_t + \alpha_{t+1} \Delta\omega_{t+1}; \tag{41}$$

$$\bar{\mathcal{J}}(\omega_t, \mathbf{z}_{t+1}) = -\left(\omega_t^{(0)} + \omega_t^{(1)} \mathbf{z}_{t+1}\right)^\top \omega_t^{(2)} \left(\omega_t^{(0)} + \omega_t^{(1)} \mathbf{z}_{t+1}\right); \tag{42}$$

► [Threshold Evaluation]

$$\gamma_{t+1} = \gamma_t - \beta_{t+1} \Delta\gamma_{t+1}(\mathbf{z}_{t+1}); \tag{43}$$

► [Tracking μ_{t+1} and Σ_{t+1} of (19) and (20)]

$$\xi_{t+1}^{(0)} = \xi_t^{(0)} + \beta_{t+1} \Delta\xi_{t+1}^{(0)}(\mathbf{z}_{t+1}); \tag{44}$$

$$\xi_{t+1}^{(1)} = \xi_t^{(1)} + \beta_{t+1} \Delta\xi_{t+1}^{(1)}(\mathbf{z}_{t+1}); \tag{45}$$

if $\theta^p \neq NULL$ **then**

$$\left.\begin{aligned} \mathbf{z}_{t+1}^p &\sim \widehat{f}_{\theta^p}(\cdot) \triangleq \lambda f_{\theta_0} + (1-\lambda) f_{\theta^p}; \\ \gamma_{t+1}^p &= \gamma_t^p - \beta_{t+1} \Delta\gamma_{t+1}(\mathbf{z}_{t+1}^p); \end{aligned}\right\} \tag{46}$$

► [Threshold Comparison]

$$T_{t+1} = T_t + c\left(I_{\{\gamma_{t+1} > \gamma_{t+1}^p\}} - I_{\{\gamma_{t+1} \leq \gamma_{t+1}^p\}} - T_t\right); \tag{47}$$

► [Updating Model Parameter]

if $T_{t+1} > \epsilon_1$ **then**

$$\theta^p = \theta_t;$$

$$\theta_{t+1} = \theta_t + \alpha_{t+1}\left((\xi_t^{(0)}, \xi_t^{(1)})^\top - \theta_t\right); \tag{48}$$

$$\gamma_{t+1}^p = \gamma_t; \quad T_t = 0; \quad c = c_t; \tag{49}$$

else

$$\gamma_{t+1}^p = \gamma_t^p; \quad \theta_{t+1} = \theta_t;$$

$t := t+1$;

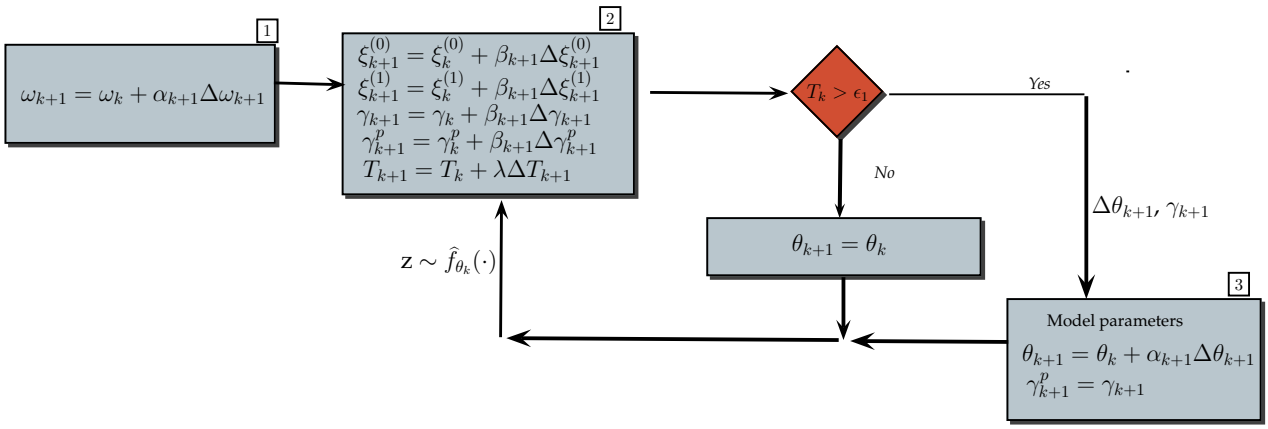

Figure 2 FlowChart representation of the algorithm SCE-MSPBEM.

REMARK 4. In practice, different stopping criteria can be used. For instance, (a) t reaches an *a priori* fixed limit, (b) the computational resources are exhausted, or (c) the variable T_t is unable to cross the ϵ_1 threshold for an *a priori* fixed number of iterations.

The pictorial depiction of the algorithm SCE-MSPBEM is shown in Figure 2.

It is important to note that the model parameter θ_t is not updated at each t. Rather it is updated every time T_t hits ϵ_1 where $0 < \epsilon_1 < 1$. So the update of θ_t only happens along a subsequence $\{t_{(n)}\}_{n\in\mathbb{Z}_+}$ of $\{t\}_{t\in\mathbb{Z}_+}$. So between $t = t_{(n)}$ and $t = t_{(n+1)}$, the variable γ_t estimates the quantity $\gamma_\rho(\bar{\mathcal{J}}_{\omega_t}, , \widehat{\theta}_{t_{(n)}})$. The threshold γ_t^p is also updated during the ϵ_1 crossover in (49). Thus $\gamma_{t_{(n)}}^p$ is the estimate of $(1-\rho)$-quantile *w.r.t.* $\widehat{f}_{\theta_{t_{(n-1)}}}$. Thus T_t in recursion (47) is a elegant trick to ensure that the estimates γ_t eventually become greater than the prior threshold $\gamma_{t_{(n)}}^p$, *i.e.*, $\gamma_t > \gamma_{t_{(n)}}^p$ for all but finitely many t. A timeline map of the algorithm is shown in Figure 3.

It can be verified as follows that the random variable T_t belongs to $(-1, 1)$, $\forall t > 0$. We state it as a proposition here.

PROPOSITION 1. *For any $T_0 \in (0,1)$, T_t in (47) belongs to $(-1,1)$, $\forall t > 0$.*

Proof: Assume $T_0 \in (0,1)$. Now the equation (47) can be rearranged as

$$T_{t+1} = (1-c)\,T_t + c(I_{\{\gamma_{t+1} > \gamma_{t+1}^p\}} - I_{\{\gamma_{t+1} \leq \gamma_{t+1}^p\}}),$$

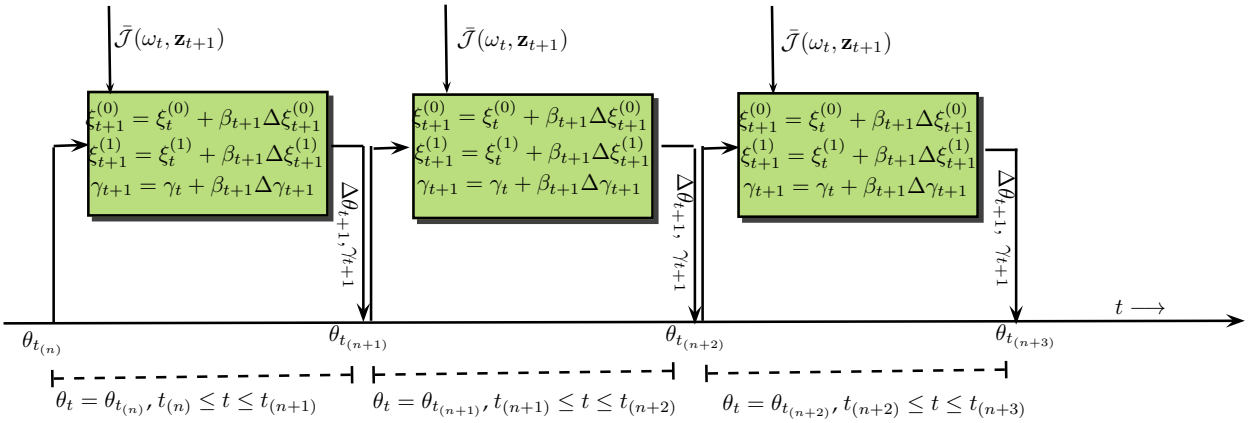

Figure 3 Timeline graph of the algorithm SCE-MSPBEM.

where $c \in (0,1)$. In the worst case, either $I_{\{\gamma_{t+1} > \gamma_{t+1}^p\}} = 1$, $\forall t$ or $I_{\{\gamma_{t+1} \le \gamma_{t+1}^p\}} = 1$, $\forall t$. Since the two events $\{\gamma_{t+1} > \gamma_{t+1}^p\}$ and $\{\gamma_{t+1} \le \gamma_{t+1}^p\}$ are mutually exclusive, we will only consider the former event $I_{\{\gamma_{t+1} > \gamma_{t+1}^p\}} = 1, \forall t$. In this case

$$\begin{aligned}\lim_{t\to\infty} T_t &= \lim_{t\to\infty} \left(c + c(1-c) + c(1-c)^2 + \cdots + c(1-c)^{t-1} + (1-c)^t T_0\right) \\ &= \lim_{t\to\infty} \frac{c(1-(1-c)^t)}{c} + T_0(1-c)^t = \lim_{t\to\infty} (1-(1-c)^t) + T_0(1-c)^t = 1. \quad (\because c \in (0,1))\end{aligned}$$

Similarly for the latter event $I_{\{\gamma_{t+1} \le \gamma_{t+1}^p\}} = 1, \forall t$, we can prove that $\lim_{t\to\infty} T_t = -1$. ■

REMARK 5. The recursion in equation (46) is not addressed in the discussion above. The update of γ_t^p in equation (49) happens along a subsequence $\{t_{(n)}\}_{n\ge 0}$. So $\gamma_{t_{(n)}}^p$ is the estimate of $\gamma_\rho(\bar{\mathcal{J}}_{\omega_{t_{(n)}}}, \theta_{t_{(n-1)}})$, where $\bar{\mathcal{J}}_{\omega_{t_{(n)}}}(\cdot) = \bar{\mathcal{J}}(\omega_{t_{(n)}}, \cdot)$. But at time $t_{(n)} < t \le t_{(n+1)}$, γ_t^p is compared with γ_t in equation (47). But γ_t is derived from a better estimate of $\bar{\mathcal{J}}(\omega_t, \cdot)$. Equation (46) ensures that γ_t^p is updated using the latest estimate of $\bar{\mathcal{J}}(\omega_t, \cdot)$. The variable θ^p holds the model parameter $\theta_{t_{(n-1)}}$ and the update of γ_t^p in (46) is performed using the $\mathbf{z}_{t+1}^p$ sampled using $\widehat{f}_{\theta^p}(\cdot)$.

3. Convergence Analysis

For analyzing the asymptotic behaviour of the algorithm, we apply the ODE based analysis from [29, 32, 28] where an ODE whose asymptotic behaviour is eventually tracked by the stochastic system is identified. The long run behaviour of the equivalent ODE is studied and it is argued that the algorithm asymptotically converges almost surely to the set of stable fixed points of the ODE. We define the filtration $\{\mathcal{F}_t\}_{t\in\mathbb{Z}_+}$ where the σ-field $\mathcal{F}_t = \sigma\left(\omega_i, \gamma_i, \gamma_i^p, \xi_i^{(0)}, \xi_i^{(1)}, \theta_i, 0 \le i \le t; \mathbf{z}_i, 1 \le i \le t; \mathbf{s}_i, \mathbf{r}_i, \mathbf{s}_i', 0 \le i < t\right)$, $t \in \mathbb{Z}_+$.

It is worth mentioning that the recursion (41) is independent of other recursions and hence can be analysed independently. For the recursion (41) we have the following result.

LEMMA 2. *Let the step-size sequences α_t and β_t , $t \in \mathbb{Z}_+$ satisfy (39). For the sample trajectory $\{(\mathbf{s}_t, \mathbf{r}_t, \mathbf{s}'_t)\}_{t=0}^{\infty}$, we let assumption (A3) hold and let ν be the sampling distribution. Then, for a given $z \in \mathcal{Z}$, the iterates ω_t in equation (41) satisfy with probability one,*

$$\lim_{t\to\infty}(\omega_t^{(0)} + \omega_t^{(1)} z) = \omega_*^{(0)} + \omega_*^{(1)} z,$$
$$\lim_{t\to\infty} \omega_t^{(2)} = \omega_*^{(2)} \text{ and } \lim_{t\to\infty} \bar{\mathcal{J}}(\omega_t, z) = \mathcal{J}(z),$$

where $\bar{\mathcal{J}}_t(z)$ is defined in equation (42), $\mathcal{J}(z)$ is defined in equation (26), Φ is defined in equation (3) and D^ν is defined in equation (4) respectively.

Proof: By rearranging equations in (41), for $t \in \mathbb{Z}_+$, we get

$$\omega_{t+1}^{(0)} = \omega_t^{(0)} + \alpha_{t+1}\big(\mathbb{M}_{t+1}^{(0,0)} + h^{(0,0)}(\omega_t^{(0)})\big), \tag{50}$$

where $\mathbb{M}_{t+1}^{(0,0)} = \mathbf{r}_t\phi_t - \mathbb{E}[\mathbf{r}_t\phi_t]$ and $h^{(0,0)}(x) = \mathbb{E}[\mathbf{r}_t\phi_t] - x$.

Similarly,

$$\omega_{t+1}^{(1)} = \omega_t^{(1)} + \alpha_{t+1}\big(\mathbb{M}_{t+1}^{(0,1)} + h^{(0,1)}(\omega_t^{(1)})\big), \tag{51}$$

where $\mathbb{M}_{t+1}^{(0,1)} = \phi_t(\gamma\phi'_t - \phi_t)^\top - \mathbb{E}[\phi_t(\gamma\phi'_t - \phi_t)^\top]$ and $h^{(0,1)}(x) = \mathbb{E}[\phi_t(\gamma\phi'_t - \phi_t)^\top] - x$.

Finally,

$$\omega_{t+1}^{(2)} = \omega_t^{(2)} + \alpha_{t+1}\big(\mathbb{M}_{t+1}^{(0,2)} + h^{(0,2)}(\omega_t^{(2)})\big), \tag{52}$$

where $\mathbb{M}_{t+1}^{(0,2)} = \mathbb{E}\left[\phi_t\phi_t^\top\omega_t^{(2)}\right] - \phi_t\phi_t^\top\omega_t^{(2)}$ and $h^{(0,2)}(x) = \mathbb{I}_{k\times k} - \mathbb{E}[\phi_t\phi_t^\top x]$.

It is easy to verify that $h^{(0,i)}, 0 \le i \le 2$ are Lipschitz continuous and $\{\mathbb{M}_{t+1}^{(0,i)}\}_{t\in\mathbb{Z}_+}$, $0 \le i \le 2$ are martingale difference noise terms, *i.e.*, for each i, $\mathbb{M}_t^{(0,i)}$ is $\mathcal{F}_t$-measurable, integrable and $\mathbb{E}\left[\mathbb{M}_{t+1}^{(0,i)}|\mathcal{F}_t\right] = 0$, $t \in \mathbb{Z}_+$, $0 \le i \le 2$.

Since ϕ_t and $\mathbf{r}_t$ have uniformly bounded second moments, the noise terms $\{\mathbb{M}_{t+1}^{(0,0)}\}_{t\in\mathbb{Z}_+}$ have uniformly bounded second moments as well and hence $\exists K_{0,0} > 0$ *s.t.*

$$\mathbb{E}\left[\|\mathbb{M}_{t+1}^{(0,0)}\|^2|\mathcal{F}_t\right] \le K_{0,0}(1 + \|\omega_t^{(0)}\|^2), t \ge 0.$$

Also $h_c^{(0,0)}(x) \triangleq \frac{h^{(0,0)}(cx)}{c} = \frac{\mathbb{E}[\mathbf{r}_t\phi_t|\mathcal{F}_t]-cx}{c} = \frac{\mathbb{E}[\mathbf{r}_t\phi_t|\mathcal{F}_t]}{c} - x$. So $h_\infty^{(0,0)}(x) = \lim_{t\to\infty} h_c^{(0,0)}(x) = -x$. Since the ODE $\dot{x}(t) = h_\infty^{(0,0)}(x)$ is globally asymptotically stable to the origin, we obtain that the iterates $\{\omega_t^{(0)}\}_{t\in\mathbb{Z}_+}$ are almost surely stable, *i.e.*, $\sup_t \|\omega_t^{(0)}\| < \infty$ *a.s.*, from Theorem 7, Chapter 3 of [28]. Similarly we can show that $\sup_t \|\omega_t^{(1)}\| < \infty$ *a.s.*

Since ϕ_t and ϕ_t' have uniformly bounded second moments, the second moments of $\{\mathbb{M}_{t+1}^{(0,1)}\}_{t\in\mathbb{Z}_+}$ are uniformly bounded and therefore $\exists K_{0,1} > 0$ *s.t.*

$$\mathbb{E}\left[\|\mathbb{M}_{t+1}^{(0,1)}\|^2|\mathcal{F}_t\right] \leq K_{0,1}(1 + \|\omega_t^{(1)}\|^2), t \geq 0.$$

Now define

$$h_c^{(0,2)}(x) \triangleq \frac{h^{(0,2)}(cx)}{c} = \frac{\mathbb{I}_{k\times k} - \mathbb{E}\left[\phi_t\phi_t^\top cx|\mathcal{F}_t\right]}{c} = \frac{\mathbb{I}_{k\times k}}{c} - x\mathbb{E}\left[\phi_t\phi_t^\top\right].$$

Hence $h_\infty^{(0,2)}(x) = \lim_{t\to\infty} h_c^{(0,2)}(x) = -x\mathbb{E}\left[\phi_t\phi_t^\top\right]$. The ∞-system ODE given by $\dot{x}(t) = h_\infty^{(0,2)}(x)$ is also globally asymptotically stable to the origin since $\mathbb{E}[\phi_t\phi_t^\top]$ is positive definite (as it is non-singular and positive semi-definite). So $\sup_t \|\omega_t^{(2)}\| < \infty$ *a.s.* from Theorem 7, Chapter 3 of [28].

Since ϕ_t has uniformly bounded second moments, $\exists K_{0,2} > 0$ *s.t.*

$$\mathbb{E}\left[\|\mathbb{M}_{t+1}^{(0,2)}\|^2|\mathcal{F}_t\right] \leq K_{0,2}(1 + \|\omega_t^{(2)}\|^2), t \geq 0.$$

Now consider the following system of ODEs associated with (50)-(52):

$$\frac{d}{dt}\omega^{(0)}(t) = \mathbb{E}\left[\mathbf{r}_t\phi_t\right] - \omega^{(0)}(t), \qquad t \in \mathbb{R}_+, \tag{53}$$

$$\frac{d}{dt}\omega^{(1)}(t) = \mathbb{E}\left[\phi_t(\gamma\phi_t' - \phi_t)^\top\right] - \omega^{(1)}(t)), \qquad t \in \mathbb{R}_+, \tag{54}$$

$$\frac{d}{dt}\omega^{(2)}(t) = \mathbb{I}_{k\times k} - \mathbb{E}\left[\phi_t\phi_t^\top\right]\omega^{(2)}(t), \qquad t \in \mathbb{R}_+. \tag{55}$$

For the ODE (53), the point $\mathbb{E}\left[\mathbf{r}_t\phi_t\right]$ is a globally asymptotically stable equilibrium. Similarly for the ODE (54), the point $\mathbb{E}\left[\phi_t(\gamma\phi_t' - \phi_t)^\top\right]$ is a globally asymptotically stable equilibrium. For the ODE (55), since $\mathbb{E}\left[\phi_t\phi_t^\top\right]$ is non-negative definite and non-singular (from the assumptions of the lemma), the ODE (55) is globally asymptotically stable to the point $\mathbb{E}\left[\phi_t\phi_t^\top\right]^{-1}$.

It can now be shown from Theorem 2, Chapter 2 of [28] that the asymptotic properties of the recursions (50), (51), (52) and their associated ODEs (53), (54), (55) are similar and hence $\lim_{t\to\infty}\omega_t^{(0)} = \mathbb{E}\left[\mathbf{r}_t\phi_t\right]$ *a.s.*, $\lim_{t\to\infty}\omega_t^{(1)} = \mathbb{E}\left[\phi_t(\gamma\phi_t' - \phi_t)^\top\right]$ *a.s.* and $\lim_{t\to\infty}\omega_{t+1}^{(2)} = \mathbb{E}\left[\phi_t\phi_t^\top\right]^{-1}$ *a.s.*. So for any $z \in \mathbb{R}^k$, using (28), we have $\lim_{t\to\infty}(\omega_t^{(0)} + \omega_t^{(1)}z) = \Phi^\top D^\nu(T_\pi V_z - V_z)$ *a.s.* Also, from (29), we have $\lim_{t\to\infty}\omega_{t+1}^{(2)} = (\Phi^\top D^\nu\Phi)^{-1}$ *a.s.*

Putting all the above together we get $\lim_{t\to\infty}\bar{\mathcal{J}}(\omega_t, z) = \bar{\mathcal{J}}(\omega_*, z) = \mathcal{J}(z)$ *a.s.* ■

As mentioned before, the update of θ_t only happens along a subsequence $\{t_{(n)}\}_{n\in\mathbb{Z}_+}$ of $\{t\}_{t\in\mathbb{Z}_+}$. So between $t = t_{(n)}$ and $t = t_{(n+1)}$, θ_t is constant. The lemma and the theorems that follow in this paper depend on the timescale difference in the step-size schedules $\{\alpha_t\}_{t\geq 0}$ and $\{\beta_t\}_{t\geq 0}$. The timescale differences allow the different recursions to learn at different rates. The step-size $\{\beta_t\}_{t\geq 0}$ decays to 0 at a slower rate than $\{\alpha_t\}_{t\geq 0}$ and hence the increments in the recursions (43), (44) and (45) which are controlled by β_t are larger and hence converge faster than the recursions (41),(42) and (48) which are controlled by α_t. So the relative evolution of the variables from the slower timescale α_t, *i.e.*, ω_t is indeed slow and in fact can be considered constant when viewed from the faster timescale β_t, see Chapter 6, [28] for a succinct description on multi-timescale stochastic approximation algorithms.

⊛ **Assumption (A4):** The iterate sequence γ_t in equation (43) satisfies $\sup_t|\gamma_t| < \infty$ *a.s.*.

REMARK 6. The assumption (A4) is a technical requirement to prove convergence. In practice, one may replace (43) by its 'projected version' whereby the iterates are projected back to an *a priori* chosen compact convex set if they stray outside of this set.

Notation: We denote by $\mathbb{E}_{\widehat{\theta}}[\cdot]$ the expectation *w.r.t.* the mixture *pdf* and $\mathbb{P}_{\widehat{\theta}}$ denotes its induced probability measure. Also $\gamma_\rho(\cdot, \widehat{\theta})$ represents the $(1-\rho)$-quantile *w.r.t.* the mixture *pdf* $\widehat{f}_\theta$.

The recursion (43) moves on a faster timescale as compared to the recursion (41) of ω_t and the recursion (48) of θ_t. Hence, on the timescale of the recursion (43), one may consider ω_t and θ_t to be fixed. For recursion (43) we have the following result:

LEMMA 3. *Let $\omega_t \equiv \omega$, $\theta_t \equiv \theta$. Let $\bar{\mathcal{J}}_\omega(\cdot) \triangleq \bar{\mathcal{J}}(\omega, \cdot)$. Then $\gamma_t, t \in \mathbb{Z}_+$ in equation (43) satisfy $\gamma_t \to \gamma_\rho(\bar{\mathcal{J}}_\omega, \widehat{\theta})$ as $t \to \infty$ with probability one.*

Proof: Here, for easy reference we rewrite the recursion (43),

$$\gamma_{t+1} = \gamma_t - \beta_{t+1} \Delta\gamma_{t+1}(\mathbf{z}_{t+1}), \tag{56}$$

Substituting the expression for $\Delta\gamma_{t+1}$ in (56) with $\omega_t = \omega$ and $\theta_t = \theta$, we get

$$\gamma_{t+1} = \gamma_t - \beta_{t+1}\Big(-(1-\rho) I_{\{\bar{\mathcal{J}}(\omega, \mathbf{z}_{t+1}) \geq \gamma_t\}} + \rho I_{\{\bar{\mathcal{J}}(\omega, \mathbf{z}_{t+1}) \leq \gamma_t\}} \Big), \quad \text{where } \mathbf{z}_{t+1} \sim \widehat{f}_\theta(\cdot) \tag{57}$$

The above equation can be apparently viewed as,

$$\gamma_{t+1} - \gamma_t \in -\beta_{t+1} \nabla_y \psi(\bar{\mathcal{J}}_\omega(\mathbf{z}_{t+1}), \gamma_t),$$

where $\nabla_y \psi$ is the sub-differential of $\psi(x, y)$ *w.r.t.* y (where ψ is defined in Lemma 1). $\nabla_y \psi$ is a set function and is defined as follows:

$$\nabla_y \psi(\bar{\mathcal{J}}_\omega(z), y) = \begin{cases} \{-(1-\rho) I_{\{\bar{\mathcal{J}}(\omega,z) \geq y\}} + \rho I_{\{\bar{\mathcal{J}}(\omega,z) \leq y\}}\}, & \text{for} \quad y \neq \bar{\mathcal{J}}(\omega, z) \\ [\rho_1, \rho_2\,], & \text{for} \quad y = \bar{\mathcal{J}}(\omega, z), \end{cases} \tag{58}$$

where $\rho_1 = \min\{1-\rho, \rho\}$ and $\rho_2 = \max\{1-\rho, \rho\}$.

Rearranging the terms in equation (57) we get,

$$\gamma_{t+1} = \gamma_t + \beta_{t+1}\left(\mathbb{M}_{t+1}^{(1,0)} - \mathbb{E}_{\widehat{\theta}}\left[\Delta\gamma_{t+1}(\mathbf{z}_{t+1})\right] \right), \tag{59}$$

where $\mathbb{M}_{t+1}^{(1,0)} = \mathbb{E}_{\widehat{\theta}}[\Delta\gamma_{t+1}(\mathbf{z}_{t+1})] - \Delta\gamma_{t+1}(\mathbf{z}_{t+1})$ with $\mathbf{z}_{t+1} \sim \widehat{f}_\theta(\cdot)$.

It is easy to verify that $\mathbb{E}_{\widehat{\theta}}[\Delta\gamma_{t+1}(\mathbf{z}_{t+1})] = \nabla_y \mathbb{E}_{\widehat{\theta}}\left[\psi(\bar{\mathcal{J}}(\mathbf{z}_{t+1}), y)\right]$. For brevity, define $h^{(1,0)}(\gamma) \triangleq -\nabla_y \mathbb{E}_{\widehat{\theta}}\left[\psi(\bar{\mathcal{J}}_\omega(\mathbf{z}_{t+1}), \gamma)\right]$.

The set function $h^{(1,0)} : \mathbb{R} \to \{\text{subsets of } \mathbb{R}\}$ satisfies the following properties:

1. For each $y \in \mathbb{R}$, $h^{(1,0)}(y)$ is convex and compact.
2. For each $y \in \mathbb{R}$, $\sup_{y' \in h(y)} |h^{(1,0)}(y')| < K_{1,0}(1 + |y|)$, for some $0 < K_{1,0} < \infty$.
3. $h^{(1,0)}$ is upper semi-continuous.

The noise term $\mathbb{M}_t^{(1,0)}$ satisfies the following properties:

1. $\mathbb{M}_t^{(1,0)}$ is $\mathcal{F}_t$-measurable $\forall t$ and integrable, $\forall t > 0$.
2. $\mathbb{M}_t^{(1,0)}$, $t \geq 0$ is a martingale difference noise sequence, *i.e.*, $\mathbb{E}[\mathbb{M}_{t+1}^{(1,0)}|\mathcal{F}_t] = 0$ *a.s.*
3. $\mathbb{E}\left[\|\mathbb{M}_{t+1}^{(1,0)}\|^2 \big| \mathcal{F}_t\right] \leq K_{1,1}(1 + \|\gamma_t\|^2 + \|\omega_t\|^2),$ for some $0 < K_{1,1} < \infty$. This follows directly from the fact that $\Delta\gamma_{t+1}(\mathbf{z}_{t+1})$ has finite first and second order moments.

Therefore by the almost sure boundedness of the sequence $\{\gamma_t\}$ in assumption (A4) and by Lemma 1, Chapter 2 in [28], we can claim that the stochastic sequence $\{\gamma_t\}$ asymptotically tracks the differential inclusion

$$\frac{d}{dt}\gamma(t) \in -\mathbb{E}_{\widehat{\theta}}\left[\nabla_y \psi(\bar{\mathcal{J}}_\omega(\mathbf{z}), \gamma(t))\right] = -\nabla_y \mathbb{E}_{\widehat{\theta}}\left[\psi(\bar{\mathcal{J}}_\omega(\mathbf{z}), \gamma(t))\right] = h^{(1,0)}(\gamma(t)). \tag{60}$$

The interchange of ∇_γ and $\mathbb{E}_{\widehat{\theta}}[\cdot]$ in the above equation is guaranteed by the dominated convergence theorem.

Now we prove the stability of the above differential inclusion. Note that by Lemma 1 of [31], we know that $\gamma^* \triangleq \gamma_\rho(\bar{\mathcal{J}}_\omega, \widehat{\theta})$ is a root of the function $h^{(1,0)}(\gamma)$ and hence it is a fixed point of the flow induced by the above differential inclusion. Now define $V(\gamma) \triangleq \mathbb{E}_{\widehat{\theta}}\left[\psi(\bar{\mathcal{J}}_\omega(\mathbf{z}), \gamma)\right] - \gamma^*$. It is easy to verify that V is continuously differentiable. Also by Lemma 1 of [31], we have $\mathbb{E}_{\widehat{\theta}}\left[\psi(\bar{\mathcal{J}}_\omega(\mathbf{z}), \gamma)\right]$ to be a convex function and γ^* to be its global minimum. Hence $V(\gamma) > 0$, $\forall \gamma \in \mathbb{R}^d \backslash \{\gamma^*\}$. Further $V(\gamma^*) = 0$ and $V(\gamma) \to \infty$ as $\|\gamma\| \to \infty$. So $V(\cdot)$ is a Lyapunov function. Also note that $< \nabla V(\gamma), h(\gamma) > \leq 0$. So $\widehat{\gamma}$ is the global attractor of the differential inclusion defined in (60). Thus by Theorem 2 of chapter 2 in [28], the iterates γ_t converge almost surely to $\gamma^* = \gamma_\rho(\omega, \widehat{\theta})$. ■

The recursions (44) and (45) move on a faster timescale as compared to the recursion (41) of ω_t and the recursion (48) of $\bar{\theta}_t$. Hence, viewed from the timescale of the recursions (44) and (45), one may consider ω_t and θ_t to be fixed. For the recursions (44) and (45), we have the following result:

LEMMA 4. *Assume $\omega_t \equiv \omega$, $\theta_t \equiv \theta$. Let $\bar{\mathcal{J}}_\omega(\cdot) \triangleq \bar{\mathcal{J}}(\omega, \cdot)$. Then almost surely,*

$$(i)\ \lim_{t\to\infty} \xi_t^{(0)} = \xi_*^{(0)} = \frac{\mathbb{E}_{\widehat{\theta}}\left[\mathbf{g}_1(\mathcal{J}_\omega(\mathbf{z}), \mathbf{z}, \gamma_\rho(\mathcal{J}, \widehat{\theta})\right]}{\mathbb{E}_{\widehat{\theta}}\left[\mathbf{g}_0(\mathcal{J}_\omega(\mathbf{z}), \gamma_\rho(\mathcal{J}, \widehat{\theta})\right]},$$

$$(ii)\ \lim_{t\to\infty} \xi_t^{(1)} = \xi_*^{(1)} = \frac{\mathbb{E}_{\widehat{\theta}}\left[\mathbf{g}_2\left(\bar{\mathcal{J}}_\omega(\mathbf{z}), \mathbf{z}, \gamma_\rho(\bar{\mathcal{J}}_\omega, \widehat{\theta}), \xi_*^{(0)}\right)\right]}{\mathbb{E}_{\widehat{\theta}}\left[\mathbf{g}_0\left(\bar{\mathcal{J}}_\omega(\mathbf{z}), \gamma_\rho(\bar{\mathcal{J}}_\omega, \widehat{\theta})\right)\right]}.$$

where $\mathbb{E}_{\widehat{\theta}}[\cdot]$ is the expectation w.r.t. the pdf $\widehat{f}_\theta(\cdot)$ and $\mathbf{z} \sim \widehat{f}_\theta(\cdot)$.

(iii) If $\gamma_\rho(\bar{\mathcal{J}}_\omega, \widehat{\theta}) > \gamma_\rho(\bar{\mathcal{J}}_\omega, \widehat{\theta^p})$, then T_t, $t \in \mathbb{Z}_+$ in equation (47) satisfy $\lim_{t\to\infty} T_t = 1$ a.s.

Proof: (i) First, we recall equation (44) below

$$\xi_{t+1}^{(0)} = \xi_t^{(0)} + \beta_{t+1}\left(\mathbf{g}_1(\bar{\mathcal{J}}(\omega_t, \mathbf{z}_{t+1}), \mathbf{z}_{t+1}, \gamma_t) - \xi_t^{(0)}\mathbf{g}_0\left(\bar{\mathcal{J}}(\omega_t, \mathbf{z}_{t+1}), \gamma_t\right)\right). \tag{61}$$

Note that the above recursion of $\xi_t^{(0)}$ depends on γ_t, but not the other way. This implies that we can replace γ_t by its limit point $\gamma_\rho(\bar{\mathcal{J}}_\omega, \widehat{\theta})$ and a bias term which goes to zero as $t \to \infty$. We denote the decaying bias term using the notation $o(1)$. Further, using the assumption that $\omega_t = \omega$, $\theta_t = \theta$ and from the equation (61), we get,

$$\xi_{t+1}^{(0)} = \xi_t^{(0)} + \beta_{t+1}\left(h^{(2,0)}(\xi_t^{(0)}) + \mathbb{M}_{t+1}^{(2,0)} + o(1)\right), \tag{62}$$

$$\text{where } h^{(2,0)}(x) \triangleq -\mathbb{E}\left[x\mathbf{g}_0\left(\bar{\mathcal{J}}_\omega(\mathbf{z}_{t+1}), \gamma_\rho(\bar{\mathcal{J}}_\omega, \widehat{\theta})\right)\Big|\mathcal{F}_t\right] + \mathbb{E}\left[\mathbf{g}_1\left(\bar{\mathcal{J}}_\omega(\mathbf{z}_{t+1}), \mathbf{z}_{t+1}, \gamma_\rho(\bar{\mathcal{J}}_\omega, \widehat{\theta})\right)\Big|\mathcal{F}_t\right], \tag{63}$$

$$\mathbb{M}_{t+1}^{(2,0)} \triangleq \mathbf{g}_1\left(\bar{\mathcal{J}}_\omega(\mathbf{z}_{t+1}), \mathbf{z}_{t+1}, \gamma_\rho(\bar{\mathcal{J}}_\omega, \widehat{\theta})\right) - \mathbb{E}\left[\mathbf{g}_1\left(\bar{\mathcal{J}}_\omega(\mathbf{z}_{t+1}), \mathbf{z}_{t+1}, \gamma_\rho(\bar{\mathcal{J}}_\omega, \widehat{\theta})\right)\Big|\mathcal{F}_t\right] -$$
$$\xi_t^{(0)}\mathbf{g}_0\left(\bar{\mathcal{J}}_\omega(\mathbf{z}_{t+1}), \gamma_\rho(\bar{\mathcal{J}}_\omega, \widehat{\theta})\right) + \mathbb{E}\left[\xi_t^{(0)}\mathbf{g}_0\left(\bar{\mathcal{J}}_\omega(\mathbf{z}_{t+1}), \gamma_\rho(\bar{\mathcal{J}}_\omega, \widehat{\theta})\right)\Big|\mathcal{F}_t\right] \text{ and } \mathbf{z}_{t+1} \sim \widehat{f}_\theta(\cdot).$$

Since $\mathbf{z}_{t+1}$ is independent of the σ-field $\mathcal{F}_t$, the function $h^{(2,0)}(\cdot)$ in equation (63) can be rewritten as

$$h^{(2,0)}(x) = -\mathbb{E}_{\widehat{\theta}}\left[x\mathbf{g}_0\left(\bar{\mathcal{J}}_\omega(\mathbf{z}), \gamma_\rho(\bar{\mathcal{J}}_\omega(\mathbf{z}), \widehat{\theta})\right)\right] + \mathbb{E}_{\widehat{\theta}}\left[\mathbf{g}_1\left(\bar{\mathcal{J}}_\omega(\mathbf{z}), \mathbf{z}, \gamma_\rho(\bar{\mathcal{J}}_\omega, \widehat{\theta})\right)\right], \text{ where } \mathbf{z} \sim \widehat{f}_\theta(\cdot).$$

It is easy to verify that $\mathbb{M}_t^{(2,0)}$, $t \in \mathbb{Z}_+$ is a martingale difference sequence, *i.e.*, $\mathbb{M}_t^{(2,0)}$ is $\mathcal{F}_t$-measurable, integrable and $\mathbb{E}[\mathbb{M}_{t+1}^{(2,0)}|\mathcal{F}_t] = 0$ *a.s.*, $\forall t \in \mathbb{Z}_+$. It is also easy to verify that $h^{(2,0)}(\cdot)$ is Lipschitz continuous. Also since $S(\cdot)$ is bounded above and $\widehat{f}_\theta(\cdot)$ has finite first and second moments we have almost surely,

$$\mathbb{E}\left[\|\mathbb{M}_{t+1}^{(2,0)}\|^2|\mathcal{F}_t\right] \leq K_{2,0}(1 + \|\xi_t^{(0)}\|^2), \forall t \geq 0, \text{ for some } 0 < K_{2,0} < \infty.$$

Now consider the ODE

$$\frac{d}{dt}\xi^{(0)}(t) = h^{(2,0)}(\xi^{(0)}(t)). \tag{64}$$

We may rewrite the above ODE as,

$$\frac{d}{dt}\xi^{(0)}(t) = A\xi^{(0)}(t) + b^{(0)},$$

where A is a diagonal matrix with $A_{ii} = -\mathbb{E}_{\widehat{\theta}}\left[\mathbf{g}_0(\bar{\mathcal{J}}_\omega(\mathbf{z}), \gamma_\rho(\bar{\mathcal{J}}_\omega, \widehat{\theta}))\right]$, $0 \leq i < k$ and $b^{(0)} = \mathbb{E}_{\widehat{\theta}}\left[\mathbf{g}_1(\bar{\mathcal{J}}_\omega(\mathbf{z}), \mathbf{z}, \gamma_\rho(\bar{\mathcal{J}}_\omega, \widehat{\theta}))\right]$. Now consider the ODE in the ∞-system $\frac{d}{dt}\xi^{(0)}(t) =$

$\lim_{\eta\to\infty}\frac{h^{(2,0)}(\eta\xi^{(0)}(t))}{\eta} = A\xi^{(0)}(t)$. Since the matrix A has the same value for all the diagonal elements, A has only one eigenvalue: $\lambda(A) = -\mathbb{E}_{\widehat{\theta}}\left[\mathbf{g}_0(\bar{\mathcal{J}}_\omega(\mathbf{z}),\gamma_\rho(\bar{\mathcal{J}}_\omega,\widehat{\theta}))\right]$ with multiplicity k. Also observe that $\lambda(A)<0$. Hence the ODE (64) is globally asymptotically stable to the origin. Using Theorem 7, Chapter 3 of [28], the iterates $\{\xi_t^{(0)}\}_{t\in\mathbb{Z}_+}$ are stable *a.s.*, *i.e.*, $\sup_{t\in\mathbb{Z}_+}\|\xi_t^{(0)}\|<\infty$ *a.s.*

Again, by using the earlier argument that the eigenvalues $\lambda(A)$ of A are negative and identical, the point $-A^{-1}b^{(0)}$ can be seen to be a globally asymptotically stable equilibrium of the ODE (64). By using Corollary 4, Chapter 2 of [28], we can conclude that

$$\lim_{t\to\infty}\xi_t^{(0)} = -A^{-1}b^{(0)}a.s. = \frac{E_{\widehat{\theta}}\left[\mathbf{g}_1(\bar{\mathcal{J}}_\omega(\mathbf{z}),\mathbf{z},\gamma_\rho(\bar{\mathcal{J}}_\omega,\widehat{\theta}))\right]}{E_{\widehat{\theta}}\left[\mathbf{g}_0(\bar{\mathcal{J}}_\omega(\mathbf{z}),\gamma_\rho(\bar{\mathcal{J}}_\omega,\widehat{\theta}))\right]}a.s.$$

(ii) We recall first the matrix recursion (45) below:

$$\xi_{t+1}^{(1)} = \xi_t^{(1)} + \beta_{t+1}\left(\mathbf{g}_2(\bar{\mathcal{J}}(\omega_t,\mathbf{z}_{t+1}),\mathbf{z}_{t+1},\gamma_t,\xi_t^{(0)}) - \xi_t^{(1)}\mathbf{g}_0\left(\bar{\mathcal{J}}(\omega_t,\mathbf{z}_{t+1}),\gamma_t\right)\right). \tag{65}$$

As in the earlier proof, we also assume $\omega_t=\omega$ and $\theta_t=\theta$. Also note that $\xi_t^{(1)}$, $\xi_t^{(0)}$ and γ_t are on the same timescale. However, the recursion of γ_t proceeds independently and in particular does not depend on $\xi_t^{(0)}$ and $\xi_t^{(1)}$. Also, there is a unilateral coupling of $\xi_t^{(1)}$ on $\xi_t^{(0)}$ and γ_t, but not the other way. Hence, while analyzing (61), one may replace γ_t and $\xi_t^{(0)}$ in equation (61) with their limit points $\gamma_\rho(\bar{\mathcal{J}}_\omega,\theta)$ and $\xi_*^{(0)}$ respectively and a decaying bias term $o(1)$. Now, by considering all the above observations, we rewrite the equation (65) as,

$$\xi_{t+1}^{(1)} = \xi_t^{(1)} + \beta_{t+1}\left(h^{(2,1)}(\xi_t^{(1)}) + \mathbb{M}_{t+1}^{(2,1)} + o(1)\right), \tag{66}$$

$$\text{where } h^{(2,1)}(x) \triangleq \mathbb{E}\left[\mathbf{g}_2\left(\bar{\mathcal{J}}_\omega(\mathbf{z}_{t+1}),\mathbf{z}_{t+1},\gamma_\rho(\bar{\mathcal{J}}_\omega,\widehat{\theta}),\xi_*^{(0)}\right)\Big|\mathcal{F}_t\right] - \mathbb{E}\left[x\mathbf{g}_0\left(\bar{\mathcal{J}}_\omega(\mathbf{z}_{t+1}),\gamma_\rho(\bar{\mathcal{J}}_\omega,\widehat{\theta})\right)\Big|\mathcal{F}_t\right] \tag{67}$$

$$\begin{aligned}\text{and } \mathbb{M}_{t+1}^{(2,1)} \triangleq\ & \mathbb{E}\left[\xi_t^{(1)}\mathbf{g}_0\left(\bar{\mathcal{J}}_\omega(\mathbf{z}_{t+1}),\gamma_\rho(\bar{\mathcal{J}}_\omega(\widehat{\theta})\right)\Big|\mathcal{F}_t\right] - \xi_t^{(1)}\mathbf{g}_0\left(\bar{\mathcal{J}}_\omega(\mathbf{z}_{t+1}),\gamma_\rho(\bar{\mathcal{J}}_\omega(\widehat{\theta})\right) - \\ & \mathbb{E}\left[\mathbf{g}_2\left(\bar{\mathcal{J}}_\omega(\mathbf{z}_{t+1}),\mathbf{z}_{t+1},\gamma_\rho(\bar{\mathcal{J}}_\omega,\widehat{\theta}),\xi_*^{(0)}\right)\Big|\mathcal{F}_t\right] + \mathbf{g}_2\left(\bar{\mathcal{J}}_\omega(\mathbf{z}_{t+1}),\mathbf{z}_{t+1},\gamma_\rho(\bar{\mathcal{J}}_\omega,\widehat{\theta}),\xi_*^{(0)}\right),\end{aligned} \tag{68}$$

where $\mathbf{z}_{t+1}\sim\widehat{f}_\theta(\cdot)$.

Since $\mathbf{z}_{t+1}$ is independent of the σ-field $\mathcal{F}_t$, the function $h^{(2,1)}(\cdot)$ in equation (67) can be rewritten as

$$h^{(2,1)}(x) = \mathbb{E}_{\widehat{\theta}}\left[\mathbf{g}_2\left(\bar{\mathcal{J}}_\omega(\mathbf{z}),\mathbf{z},\gamma_\rho(\bar{\mathcal{J}}_\omega,\widehat{\theta}),\xi_*^{(0)}\right)\right] - \mathbb{E}_{\widehat{\theta}}\left[x\mathbf{g}_0\left(\bar{\mathcal{J}}_\omega(\mathbf{z}),\gamma_\rho(\bar{\mathcal{J}}_\omega,\widehat{\theta})\right)\right],\ \text{where } \mathbf{z}\sim\widehat{f}_\theta(\cdot). \tag{69}$$

It is not difficult to verify that $\mathbb{M}_{t+1}^{(2,1)}$, $t \in \mathbb{Z}_+$ is a martingale difference noise sequence and $h^{(2,1)}(\cdot)$ is Lipschitz continuous. Also since $S(\cdot)$ is bounded and $\widehat{f}_\theta(\cdot)$ has finite first and second moments we get,

$$\mathbb{E}\left[\|\mathbb{M}_{t+1}^{(2,1)}\|^2|\mathcal{F}_t\right] \leq K_{2,1}(1+\|\xi_t^{(1)}\|^2), \forall t \in \mathbb{Z}_+, \text{ for some } 0 < K_{2,1} < \infty.$$

Now consider the ODE given by

$$\frac{d}{dt}\xi^{(1)}(t) = h^{(2,1)}(\xi^{(1)}(t)), \quad t \in \mathbb{R}_+. \tag{70}$$

By rewriting the above equation we get,

$$\frac{d}{dt}\xi^{(1)}(t) = A\xi^{(1)}(t) + b^{(1)}, \quad t \in \mathbb{R}_+,$$

where A is a diagonal matrix as before, *i.e.*, $A_{ii} = -\mathbb{E}_{\widehat{\theta}}\left[\mathbf{g}_0(\bar{\mathcal{J}}_\omega(\mathbf{z}), \gamma_\rho(\bar{\mathcal{J}}_\omega, \widehat{\theta})\right]$, $\forall i, 0 \leq i < k$ and $b^{(1)} = \mathbb{E}_{\widehat{\theta}}\left[\mathbf{g}_2(\bar{\mathcal{J}}_\omega(\mathbf{z}), \mathbf{z}, \gamma_\rho(\bar{\mathcal{J}}_\omega, \widehat{\theta}), \xi_*^{(0)})\right]$. Now consider the ODE in the ∞-system $\frac{d}{dt}\xi^{(1)}(t) = \lim_{\eta\to\infty}\frac{1}{\eta}h^{(2,1)}(\eta\xi^{(1)}(t)) = A\xi^{(1)}(t)$. Again, the eigenvalue $\lambda(A) = -\mathbb{E}_{\widehat{\theta}}\left[\mathbf{g}_0(\bar{\mathcal{J}}_\omega(\mathbf{z}), \gamma_\rho(\bar{\mathcal{J}}_\omega, \widehat{\theta})\right]$ of A is negative and is of multiplicity k and hence origin is the unique globally asymptotically stable equilibrium of the ∞-system. Therefore it follows that the iterates $\{\xi_t^{(1)}\}_{t\in\mathbb{Z}_+}$ are almost surely stable, *i.e.*, $\sup_{t\in\mathbb{Z}_+}\|\xi_t^{(0)}\| < \infty$ *a.s.*, see Theorem 7, Chapter 3 of [28].

Again, by using the earlier argument that the eigenvalues $\lambda(A)$ of A are negative and identical, the point $-A^{-1}b^{(1)}$ can be seen to be a globally asymptotically stable equilibrium of the ODE (70). By Corollary 4, Chapter 2 of [28], it follows that

$$\lim_{t\to\infty}\xi_t^{(1)} = -A^{-1}b^{(1)} \ \ a.s. = \frac{\mathbb{E}_{\widehat{\theta}}\left[\mathbf{g}_2(\bar{\mathcal{J}}_\omega(\mathbf{z}), \mathbf{z}, \gamma_\rho(\bar{\mathcal{J}}_\omega, \widehat{\theta}), \xi_*^{(0)})\right]}{\mathbb{E}_{\widehat{\theta}}\left[\mathbf{g}_0(\bar{\mathcal{J}}_\omega(\mathbf{z}), \gamma_\rho(\bar{\mathcal{J}}_\omega, \widehat{\theta})\right]} \ \ a.s.$$

(iii) Here also we assume $\omega_t = \omega$. Then γ_t in recursion (43) and γ_t^p in recursion (46) converge to $\gamma_\rho(\bar{\mathcal{J}}_\omega, \widehat{\theta})$ and $\gamma_\rho(\bar{\mathcal{J}}_\omega, \widehat{\theta^p})$ respectively. So if $\gamma_\rho(\bar{\mathcal{J}}_\omega, \widehat{\theta}) > \gamma_\rho(\bar{\mathcal{J}}_\omega, \widehat{\theta^p})$, then $\gamma_t > \gamma_t^p$ eventually, *i.e.*, $\gamma_t > \gamma_t^p$ for all but finitely many t. So almost surely T_t in equation (47) will converge to $\mathbb{E}\left[I_{\{\gamma_t > \gamma_t^p\}} - I_{\{\gamma_t \leq \gamma_t^p\}}\right] = \mathbb{P}\{\gamma_t > \gamma_t^p\} - \mathbb{P}\{\gamma_t \leq \gamma_t^p\} = 1 - 0 = 1$. ■

Notation: For the subsequence $\{t_{(n)}\}_{n>0}$ of $\{t\}_{t\geq 0}$, we denote $t_{(n)}^- \triangleq t_{(n)} - 1$ for $n > 0$.

As mentioned earlier, $\bar{\theta}_t$ is updated only along a subsequence $\{t_{(n)}\}_{n\geq 0}$ of $\{t\}_{t\geq 0}$ with $t_0 = 0$ as follows:

$$\bar{\theta}_{t_{(n+1)}} = \bar{\theta}_{t_{(n)}} + \alpha_{t_{(n+1)}}\left((\xi_{t_{(n+1)}^-}^{(0)}, \xi_{t_{(n+1)}^-}^{(1)})^\top - \bar{\theta}_{t_{(n)}}\right). \tag{71}$$

Now define $\Psi(\omega,\theta) = (\Psi_1(\omega,\theta), \Psi_2(\omega,\theta))^\top$, where

$$\Psi_1(\omega,\theta) \triangleq \frac{\mathbb{E}_{\widehat{\theta}}\left[\mathbf{g}_1(\bar{\mathcal{J}}_\omega(\mathbf{z}), \mathbf{z}, \gamma_\rho(\bar{\mathcal{J}}_\omega, \widehat{\theta})\right]}{\mathbb{E}_{\widehat{\theta}}\left[\mathbf{g}_0(\bar{\mathcal{J}}_\omega(\mathbf{z}), \gamma_\rho(\bar{\mathcal{J}}_\omega, \widehat{\theta})\right]}, \tag{72}$$

$$\Psi_2(\omega,\theta) \triangleq \frac{\mathbb{E}_{\widehat{\theta}}\left[\mathbf{g}_2\left(\bar{\mathcal{J}}_\omega(\mathbf{z}), \mathbf{z}, \gamma_\rho(\bar{\mathcal{J}}_\omega, \widehat{\theta}), \Psi_1(\omega,\theta)\right)\right]}{\mathbb{E}_{\widehat{\theta}}\left[\mathbf{g}_0\left(\bar{\mathcal{J}}_\omega(\mathbf{z}), \gamma_\rho(\bar{\mathcal{J}}_\omega, \widehat{\theta})\right)\right]}. \tag{73}$$

We now state our main theorem. The theorem states that the model sequence $\{\theta_t\}$ generated by Algorithm 2 converges to $\theta^* = (z^*, 0_{k\times k})^\top$, which is the degenerate distribution concentrated at z^*.

THEOREM 2. *Let $S(z) = exp(rz)$, $r \in \mathbb{R}_+$. Let $\rho \in (0,1)$, $\lambda \in (0,1)$ and $\lambda > \rho$. Let $\theta_0 = (\mu_0, qI_{k\times k})^\top$, where $q \in \mathbb{R}_+$. Let the step-size sequences α_t, β_t, $t \in \mathbb{Z}_+$ satisfy (39). Also let $c_t \to 0$. Suppose $\{\theta_t = (\mu_t, \Sigma_t)^\top\}_{t\in\mathbb{Z}_+}$ is the sequence generated by Algorithm 2 and assume $\theta_t \in int(\Theta)$, $\forall t \in \mathbb{Z}_+$. Also, let the assumptions (A1), (A2), (A3) and (A4) hold. Further, we assume that there exists a continuously differentiable function $V : \Theta \to \mathbb{R}_+$ s.t. $\nabla V^\top(\theta)\Psi(\omega_*,\theta) < 0$, $\forall \theta \in \Theta \setminus \{\theta^*\}$ and $\nabla V^\top(\theta^*)\Psi(\omega_*,\theta^*) = 0$. Then, there exists $q^* \in \mathbb{R}_+$, $r^* \in \mathbb{R}_+$ and $\rho^* \in (0,1)$ s.t. $\forall q > q^*$, $\forall r > r^*$ and $\forall \rho < \rho^*$,*

$$\lim_{t\to\infty} \bar{\mathcal{J}}(\omega_t, \mu_t) = \mathcal{J}^* \quad \textit{and} \quad \lim_{t\to\infty} \theta_t = \theta^* = (z^*, 0_{k\times k})^\top \textit{ almost surely},$$

where $\mathcal{J}^*$ and z^* are defined in (26). Further, since $\mathcal{J} = -\mathrm{MSPBE}$, the algorithm SCE-MSPBEM converges to the global minimum of MSPBE a.s.

Proof: Rewriting the equation (48) along the subsequence $\{t_{(n)}\}_{n\in\mathbb{Z}_+}$, we have for $n \in \mathbb{Z}_+$,

$$\theta_{t_{(n+1)}} = \theta_{t_{(n)}} + \alpha_{t_{(n+1)}}\left((\xi^{(0)}_{t^-_{(n+1)}}, \xi^{(1)}_{t^-_{(n+1)}})^\top - \theta_{t_{(n)}}\right). \tag{74}$$

The iterates $\theta_{t_{(n)}}$ are stable, *i.e.*, $\sup_n \|\theta_{t_{(n)}}\| < \infty$ *a.s.* It is directly implied from the assumptions that $\theta_{t_{(n)}} \in int(\Theta)$ and Θ is a compact set.

Rearranging the equation (74) we get, for $n \in \mathbb{Z}_+$,

$$\theta_{t_{(n+1)}} = \theta_{t_{(n)}} + \alpha_{t_{(n+1)}}\left(\Psi(\omega_*, \theta_{t_{(n)}}) + o(1)\right). \tag{75}$$

This easily follows from the fact that, for $t_{(n)} < t \le t_{(n+1)}$, the random variables $\xi^{(0)}_t$ and $\xi^{(1)}_t$ estimates the quantities $\Psi_1(\omega_{t_{(n)}}, \theta_{t_{(n)}})$ and $\Psi_2(\omega_{t_{(n)}}, \theta_{t_{(n)}})$ respectively. Since $c_t \to 0$, the estimation error decays to 0.

Hence the term $o(1)$.

The limit points of the above recursion are the roots of Ψ. Hence by equating $\Psi_1(\omega_*, \theta)$ to $0_{k\times 1}$, we get,

$$\mu = \frac{\mathbb{E}_{\widehat{\theta}}\left[\mathbf{g_1}\big(\mathcal{J}(\mathbf{z}), \mathbf{z}, \gamma_\rho(\mathcal{J}, \widehat{\theta})\big)\right]}{\mathbb{E}_{\widehat{\theta}}\left[\mathbf{g_0}\big(\mathcal{J}(\mathbf{z}), \gamma_\rho(\mathcal{J}, \widehat{\theta})\big)\right]}. \tag{76}$$

Equating $\Psi_2(\omega_*, \theta)$ to $\mathbb{O}$ $(= 0_{k\times k})$, we get,

$$\frac{\mathbb{E}_{\widehat{\theta}}\left[\mathbf{g_2}\big(\mathcal{J}(\mathbf{z}), \mathbf{z}, \gamma_\rho(\mathcal{J}, \widehat{\theta}), \mu\big)\right]}{\mathbb{E}_{\widehat{\theta}}\left[\mathbf{g_0}\big(\mathcal{J}(\mathbf{z}), \gamma_\rho(\mathcal{J}, \widehat{\theta})\big)\right]} - \Sigma = \mathbb{O}. \tag{77}$$

For brevity, we define

$$\gamma_\rho^*(\theta) \triangleq \gamma_\rho(\mathcal{J}, \theta), \quad \hat{\mathbf{g}}_\mathbf{0}(z, \theta) \triangleq \mathbf{g_0}\big(\mathcal{J}(z), \gamma_\rho^*(\theta)\big) \text{ and } L(\theta) \triangleq \mathbb{E}_\theta\left[\hat{\mathbf{g}}_\mathbf{0}(\mathbf{z}, \theta)\right]. \tag{78}$$

Substituting the expression for μ from (76) in (77) and after further simplification we get,

$$(1/L(\widehat{\theta}))\mathbb{E}_{\widehat{\theta}}\left[\hat{\mathbf{g}}_\mathbf{0}(\mathbf{z}, \widehat{\theta})\mathbf{z}\mathbf{z}^\top\right] - \mu\mu^\top - \Sigma = \mathbb{O}.$$

Since $\Sigma = \mathbb{E}_\theta\left[\mathbf{z}\mathbf{z}^\top\right] - \mu\mu^\top$, the above equation implies

$$\begin{aligned}
&(1/L(\widehat{\theta}))\mathbb{E}_{\widehat{\theta}}\left[\hat{\mathbf{g}}_\mathbf{0}(\mathbf{z}, \widehat{\theta})\mathbf{z}\mathbf{z}^\top\right] - \mathbb{E}_\theta\left[\mathbf{z}\mathbf{z}^\top\right] = \mathbb{O} \\
\implies_1 \quad &\mathbb{E}_{\widehat{\theta}}\left[\hat{\mathbf{g}}_\mathbf{0}(\mathbf{z}, \widehat{\theta})\mathbf{z}\mathbf{z}^\top\right] - L(\widehat{\theta})\mathbb{E}_\theta\left[\mathbf{z}\mathbf{z}^\top\right] = \mathbb{O}, \\
\implies_2 \quad &(1-\lambda)\mathbb{E}_\theta\left[\hat{\mathbf{g}}_\mathbf{0}(\mathbf{z}, \widehat{\theta})\mathbf{z}\mathbf{z}^\top\right] + \lambda\mathbb{E}_{\theta_0}\left[\hat{\mathbf{g}}_\mathbf{0}(\mathbf{z}, \widehat{\theta})\mathbf{z}\mathbf{z}^\top\right] - L(\widehat{\theta})\mathbb{E}_\theta\left[\mathbf{z}\mathbf{z}^\top\right] = \mathbb{O} \\
\implies_3 \quad &(1-\lambda)\mathbb{E}_\theta\left[\hat{\mathbf{g}}_\mathbf{0}(\mathbf{z}, \widehat{\theta})\mathbf{z}\mathbf{z}^\top\right] + \lambda\mathbb{E}_{\theta_0}\left[\hat{\mathbf{g}}_\mathbf{0}(\mathbf{z}, \widehat{\theta})\mathbf{z}\mathbf{z}^\top\right] - (1-\lambda)\mathbb{E}_\theta\left[\hat{\mathbf{g}}_\mathbf{0}(\mathbf{z}, \widehat{\theta})\right]\mathbb{E}_\theta\left[\mathbf{z}\mathbf{z}^\top\right] - \\
&\qquad\qquad \lambda\mathbb{E}_{\theta_0}\left[\hat{\mathbf{g}}_\mathbf{0}(\mathbf{z}, \widehat{\theta})\right]\mathbb{E}_\theta\left[\mathbf{z}\mathbf{z}^\top\right] = \mathbb{O} \\
\implies_4 \quad &(1-\lambda)\mathbb{E}_\theta\left[\left(\hat{\mathbf{g}}_\mathbf{0}(\mathbf{z}, \widehat{\theta}) - \mathbb{E}_\theta\left[\hat{\mathbf{g}}_\mathbf{0}(\mathbf{z}, \widehat{\theta})\right]\right)\mathbf{z}\mathbf{z}^\top\right] + \lambda\mathbb{E}_{\theta_0}\left[\left(\hat{\mathbf{g}}_\mathbf{0}(\mathbf{z}, \widehat{\theta}) - \mathbb{E}_{\theta_0}\left[\hat{\mathbf{g}}_\mathbf{0}(\mathbf{z}, \widehat{\theta})\right]\right)\mathbf{z}\mathbf{z}^\top\right] + \\
&\qquad\qquad - \lambda\left(\mathbb{E}_\theta\left[\mathbf{z}\mathbf{z}^\top\right] - \mathbb{E}_{\theta_0}\left[\mathbf{z}\mathbf{z}^\top\right]\right)\mathbb{E}_{\theta_0}\left[\hat{\mathbf{g}}_\mathbf{0}(\mathbf{z}, \widehat{\theta})\right] = \mathbb{O} \\
\implies_5 \quad &(1-\lambda)\Sigma^2\mathbb{E}_\theta\left[\nabla_z^2\hat{\mathbf{g}}_\mathbf{0}(\mathbf{z}, \widehat{\theta})\right] + \lambda\Sigma_0^2\mathbb{E}_{\theta_0}\left[\nabla_z^2\mathbf{g}_\mathbf{0}(\mathbf{z}, \widehat{\theta})\right] + \\
&\qquad\qquad \lambda\left(\mathbb{E}_\theta\left[\mathbf{z}\mathbf{z}^\top\right] - \mathbb{E}_{\theta_0}\left[\mathbf{z}\mathbf{z}^\top\right]\right)\mathbb{E}_{\theta_0}\left[\hat{\mathbf{g}}_\mathbf{0}(\mathbf{z}, \widehat{\theta})\right] = \mathbb{O} \\
\implies_6 \quad &(1-\lambda)\Sigma^2\mathbb{E}_\theta\left[\nabla_z^2\hat{\mathbf{g}}_\mathbf{0}(\mathbf{z}, \widehat{\theta})\right] + \lambda q^2\mathbb{E}_{\theta_0}\left[\nabla_z^2\hat{\mathbf{g}}_\mathbf{0}(\mathbf{z}, \widehat{\theta})\right] + \\
&\qquad\qquad \lambda\left(\mathbb{E}_\theta\left[\mathbf{z}\mathbf{z}^\top\right] - \mathbb{E}_{\theta_0}\left[\mathbf{z}\mathbf{z}^\top\right]\right)\mathbb{E}_{\theta_0}\left[\hat{\mathbf{g}}_\mathbf{0}(\mathbf{z}, \widehat{\theta})\right] = \mathbb{O}
\end{aligned}$$

$$\begin{aligned}\implies_8 \quad (1-\lambda)\Sigma^2\mathbb{E}_\theta\left[S(\mathcal{J}(\mathbf{z}))G^r(\mathbf{z})I_{\{\mathcal{J}(\mathbf{z})\geq\gamma^*_\rho(\widehat{\theta})\}}\right]+\lambda q^2\mathbb{E}_{\theta_0}\left[S(\mathcal{J}(\mathbf{z}))G^r(\mathbf{z})I_{\{\mathcal{J}(\mathbf{z})\geq\gamma^*_\rho(\widehat{\theta})\}}\right]+\\ \lambda\left(\mathbb{E}_\theta\left[\mathbf{z}\mathbf{z}^\top\right]-\mathbb{E}_{\theta_0}\left[\mathbf{z}\mathbf{z}^\top\right]\right)\mathbb{E}_{\theta_0}\left[\hat{\mathbf{g}}_\mathbf{0}(\mathbf{z},\widehat{\theta})\right]=\mathbb{O},\end{aligned} \tag{79}$$

where $G^r(z)\triangleq r^2\nabla\mathcal{J}(z)\nabla\mathcal{J}(z)^\top+r\nabla^2\mathcal{J}(z)$. Note that $\implies_5$ follows from "integration by parts" rule for multivariate Gaussian and $\implies_8$ follows from the assumption $S(z)=exp(rz)$. Note that for each $z\in\mathcal{Z}$, $G^r(z)\in\mathbb{R}^{k\times k}$. Hence we denote $G^r(z)$ as $\left[G^r_{ij}(z)\right]_{i=1,j=1}^{i=k,j=k}$. For brevity, we also define

$$F^{r,\rho}(z,\theta)\triangleq S(\mathcal{J}(z))G^r(z)I_{\{\mathcal{J}(z)\geq\gamma^*_\rho(\theta)\}}, \tag{80}$$

where $F^{r,\rho}(z,\theta)\in\mathbb{R}^{k\times k}$ which is also denoted as $\left[F^{r,\rho}_{ij}(z)\right]_{i=1,j=1}^{i=k,j=k}$.

Hence equation (79) becomes,

$$(1-\lambda)\Sigma^2\mathbb{E}_\theta\left[F^{r,\rho}(\mathbf{z},\widehat{\theta})\right]+\lambda q^2\mathbb{E}_{\theta_0}\left[F^{r,\rho}(\mathbf{z},\widehat{\theta})\right]+\lambda\left(\mathbb{E}_\theta\left[\mathbf{z}\mathbf{z}^\top\right]-\mathbb{E}_{\theta_0}\left[\mathbf{z}\mathbf{z}^\top\right]\right)\mathbb{E}_{\theta_0}\left[\hat{\mathbf{g}}_\mathbf{0}(\mathbf{z},\widehat{\theta})\right]=\mathbb{O}. \tag{81}$$

Note that $(\nabla_i\mathcal{J})^2\geq 0$. Hence we can find a $r^*\in\mathbb{R}_+$ *s.t.* $G^r_{ii}(z)>0$, $\forall r>r^*$, $1\leq i\leq k$, $\forall z\in\mathcal{Z}$. This further implies that $\mathbb{E}_\theta[F^{r,\rho}_{ii}(\mathbf{z},\widehat{\theta}]>0$, $\forall\theta\in\Theta$. Also since $\mathcal{Z}$ is compact and $\mathcal{J}$ is continuous, we have $J(z)>B_1>-\infty$, $\forall z\in\mathcal{Z}$. Hence we obtain the following bound:

$$\left(\mathbb{E}_\theta\left[\mathbf{z}_i^2\right]-\mathbb{E}_{\theta_0}\left[\mathbf{z}_i^2\right]\right)\mathbb{E}_{\theta_0}\left[\hat{\mathbf{g}}_\mathbf{0}(\mathbf{z},\widehat{\theta})\right]>\rho K_3 S(B_1),\text{ where } 0<K_3<\infty. \tag{82}$$

Now from (81) and (82), we can find a $\rho^*\in(0,1)$ and $q^*\in\mathbb{R}$ *s.t.* $\forall\rho<\rho^*$, $\forall q>q^*$, $\forall r>r^*$, we have,

$$(1-\lambda)\Sigma^2\mathbb{E}_\theta\left[F^{r,\rho}_{ii}(\mathbf{z},\widehat{\theta})\right]+\lambda q^2\mathbb{E}_{\theta_0}\left[F^{r,\rho}_{ii}(\mathbf{z},\widehat{\theta})\right]+\lambda\left(\mathbb{E}_\theta\left[\mathbf{z}_i^2\right]-\mathbb{E}_{\theta_0}\left[\mathbf{z}_i^2\right]\right)\mathbb{E}_{\theta_0}\left[\hat{\mathbf{g}}_\mathbf{0}(\mathbf{z},\widehat{\theta})\right]>0. \tag{83}$$

This contradicts equation (81) for such a choice of ρ, q and r. This implies that each of the terms in equation (81) is 0, *i.e.*,

$$\Sigma^2\mathbb{E}_\theta\left[F^{r,\rho}(\mathbf{z},\widehat{\theta})\right]=\mathbb{O}, \tag{84}$$

$$q^2\mathbb{E}_{\theta_0}\left[F^{r,\rho}(\mathbf{z},\widehat{\theta})\right]=\mathbb{O}\qquad\text{and} \tag{85}$$

$$\left(\mathbb{E}_\theta\left[\mathbf{z}\mathbf{z}^\top\right]-\mathbb{E}_{\theta_0}\left[\mathbf{z}\mathbf{z}^\top\right]\right)\mathbb{E}_{\theta_0}\left[\hat{\mathbf{g}}_\mathbf{0}(\mathbf{z},\widehat{\theta})\right]=\mathbb{O}. \tag{86}$$

It is easy to verify that $\Sigma=\mathbb{O}$ simultaneously satisfies (84), (85) and (86). Besides $\Sigma=\mathbb{O}$ is the only solution for $r>r^*$. This is because we have already established earlier that $\forall r>r^*$, $1\leq i\leq k$, $\forall z\in\mathcal{Z}$

$\mathbb{E}_\theta[F_{ii}^{r,\rho}(\mathbf{z},\widehat{\theta}] > 0$, $\forall \theta \in \Theta$. This proves that for any $z \in \mathcal{Z}$, the degenerate distribution concentrated on z given by $\theta_z = (z, 0_{k\times k})^\top$ is a potential limit point of the recursion. From (85), we have

$$\mathbb{E}_{\theta_0}\left[F^{r,\rho}(\mathbf{z},\widehat{\theta})\right] = \mathbb{O} \qquad \Longrightarrow \qquad \gamma_\rho^*(\widehat{\theta}) = \mathcal{J}(z^*).$$

Claim A: The only degenerate distribution which satisfies the above condition is $\theta^* = (z^*, 0_{k\times k})^\top$.

The above claim can be verified as follows: if there exists $z'(\in \mathcal{Z}) \neq z^*$ *s.t.* $\gamma_\rho^*(\widehat{\theta}_{z'}) = \mathcal{J}(z^*)$ is satisfied (where $\widehat{\theta}_{z'}$ represents the mixture distribution $\widehat{f}_{\theta_{z'}}$) , then from the definition of $\gamma_\rho^*(\cdot)$ in (16) and (78), we can find an increasing sequence $\{l_i\}$, where $l_i > J(z')$ *s.t.* the following property is satisfied:

$$\lim_{i\to\infty} l_i = \mathcal{J}(z^*) \text{ and } \mathbb{P}_{\widehat{\theta}}(\mathcal{J}(z) \geq l_i) \geq \rho. \tag{87}$$

But $\mathbb{P}_{\widehat{\theta_{z'}}}(\mathcal{J}(z) \geq l_i) = (1-\lambda)\mathbb{P}_{\theta_{z'}}(\mathcal{J}(z) \geq l_i) + \lambda\mathbb{P}_{\theta_0}(\mathcal{J}(z) \geq l_i)$ and $\mathbb{P}_{\theta_{z'}}(\mathcal{J}(z) \geq l_i) = 0$, $\forall i$. Therefore from (87), we get,

$$\begin{aligned}
&\mathbb{P}_{\widehat{\theta_{z'}}}(\mathcal{J}(z) \geq l_i) \geq \rho \\
&\Rightarrow (1-\lambda)\mathbb{P}_{\theta_{z'}}(\mathcal{J}(z) \geq l_i) + \lambda\mathbb{P}_{\theta_0}(\mathcal{J}(z) \geq l_i) \geq \rho \\
&\Rightarrow \lambda\mathbb{P}_{\theta_0}(\mathcal{J}(z) \geq l_i) \geq \rho \\
&\Rightarrow \mathbb{P}_{\theta_0}(\mathcal{J}(z) \geq l_i) \geq \frac{\rho}{\lambda} < 1.
\end{aligned}$$

Recollect that $l_i \to \mathcal{J}(z^*)$. Thus by the continuity of probability measures we get

$$0 = \mathbb{P}_{\theta_0}(\mathcal{J}(z) \geq J(z^*)) = \lim_{i\to\infty} \mathbb{P}_{\theta_0}(\mathcal{J}(z) \geq l_i) \geq \frac{\rho}{\lambda},$$

which is a contradiction. This proves the Claim A. Now the only remaining task is to prove θ^* is a stable attractor. This easily follows from the assumption regarding the existence of the Lyapunov function V in the statement of the theorem. ■

3.1. Computational Complexity

The computational load of this algorithm is $\Theta(k^2)$ *per iteration* which comes from (41). Least squares algorithms like LSTD and LSPE also require $\Theta(k^2)$ per iteration. However, LSTD requires an extra operation of inverting the $k \times k$ matrix A_T which requires an extra computational effort of $\Theta(k^3)$. (Note that LSPE

also requires a $k \times k$ matrix inversion). This makes the *overall complexity of LSTD and LSPE to be* $\Theta(k^3)$. Further in some cases the matrix A_T may not be invertible. In that case, the pseudo inverse of A_T needs to be obtained in LSTD, LSPE which is computationally even more expensive. Our algorithm does not require such an inversion procedure. Also *even though the complexity of the first order temporal difference algorithms such as TD(λ) and GTD2 is* $\Theta(k)$*, the approximations they produced in the experiments we conducted turned out to be inferior to ours and also showed a slower rate of convergence than our algorithm.* Another noteworthy characteristic exhibited by our algorithm is *stability*. Recall that the convergence of TD(0) is guaranteed by the requirements that the Markov Chain of P^{π} should be ergodic and the sampling distribution ν to be its stationary distribution. The classic example of Baird's 7-star [10] violates those restrictions and hence TD(0) is seen to diverge. However, our algorithm does not impose such restrictions and shows stable behaviour even in non-ergodic off policy cases such as the Baird's example.

4. Experimental Results

We present here a numerical comparison of SCE-MSPBEM with various state-of-the-art algorithms in the literature on some benchmark Reinforcement Learning problems. In each of the experiments, a random trajectory $\{(s_t, r_t, s'_t)\}_{t=0}^{\infty}$ is chosen and all the algorithms are updated using it. Each s_t in $\{(s_t, r_t, s'_t), t \geq 0\}$ is sampled using an arbitrary distribution ν over $\mathbb{S}$. The algorithms are run on multiple trajectories and the average of the results obtained are plotted. The x-axis in the plots is $t/1000$, where t is the iteration number. In each case, the learning rates α_t, β_t are chosen so that the condition (39) is satisfied. The function $S(\cdot)$ is chosen as $S(x) = \exp(rx)$, where $r \in \mathbb{R}$ is chosen appropriately.

SCE-MSPBEM was tested on the following benchmark problems:

1. Linearized Cart-Pole Balancing [14]
2. 5-Link Actuated Pendulum Balancing [14]
3. Baird's 7-Star MDP [10]
4. 10-state Ring MDP [33]
5. Large state space and action space with Radial Basis Functions
6. Large state space and action space with Fourier Basis Functions [5]

4.1. Experiment 1: Linearized Cart-Pole Balancing [14]

Setup: A pole with mass m and length l is connected to a cart of mass M . It can rotate 360° and the cart is free to move in either direction within the bounds of a linear track.

Goal: To balance the pole upright and the cart at the centre of the track.

State space: The 4-tuple $[x, \dot{x}, \psi, \dot{\psi}]$ where ψ is the angle of the pendulum *w.r.t.* the vertical axis, $\dot{\psi}$ is the angular velocity, x the relative cart position from the centre of the track and $\dot{x}$ is its velocity.

Control space: The controller applies a horizontal force a on the cart parallel to the track. The stochastic policy used in this setting corresponds to $\pi(a|s) = \mathcal{N}(a|\beta_1^\top s, \sigma_1^2)$.

System dynamics: The dynamical equations of the system are given by

$$\ddot{\psi} = \frac{-3ml\dot{\psi}^2 \sin\psi\cos\psi + (6M+m)g\sin\psi - 6(a - b\dot{\psi})\cos\psi}{4l(M+m) - 3ml\cos\psi}, \tag{88}$$

$$\ddot{x} = \frac{-2ml\dot{\psi}^2 \sin\psi + 3mg\sin\psi\cos\psi + 4a - 4b\dot{\psi}}{4(M+m) - 3m\cos\psi}. \tag{89}$$

By making further assumptions on the initial conditions, the system dynamics can be approximated accurately by the linear system

$$\begin{bmatrix} x_{t+1} \\ \dot{x}_{t+1} \\ \psi_{t+1} \\ \dot{\psi}_{t+1} \end{bmatrix} = \begin{bmatrix} x_t \\ \dot{x}_t \\ \psi_t \\ \dot{\psi}_t \end{bmatrix} + \Delta t \begin{bmatrix} \dot{\psi}_t \\ \frac{3(M+m)\psi_t - 3a + 3b\dot{\psi}_t}{4Ml - ml} \\ \dot{x}_t \\ \frac{3mg\psi_t + 4a - 4b\dot{\psi}_t}{4M - m} \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ 0 \\ \mathbf{z} \end{bmatrix}, \tag{90}$$

where Δt is the integration time step, *i.e.*, the time difference between two transitions and $\mathbf{z}$ is a Gaussian noise on the velocity of the cart with standard deviation σ_2.

Reward function: $\mathrm{R}(s,a) = \mathrm{R}(\psi, \dot{\psi}, x, \dot{x}, a) = -100\psi^2 - x^2 - \frac{1}{10}a^2$.

Feature vectors: $\phi(s \in \mathbb{R}^4) = (1, s_1^2, s_2^2 \ldots, s_1 s_2, s_1 s_3, \ldots, s_3 s_4)^\top \in \mathbb{R}^{11}$.

Evaluation policy: The policy evaluated in the experiment is the optimal policy $\pi^*(a|s) = \mathcal{N}(a|{\beta_1^*}^\top s, {\sigma_1^*}^2)$. The parameters β_1^* and σ_1^* are computed using dynamic programming. The feature set chosen above is a perfect feature set, *i.e.*, $V^{\pi^*} \in \{\Phi z | z \in \mathbb{R}^k\}$. The table of the various parameter values we used in our

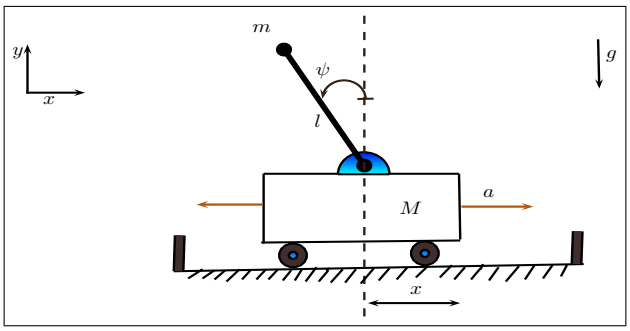

Figure 4 The Cart-Pole System. The goal is to keep the pole in the upright position and the cart at the center of the track by pushing the cart with a force a either to the left or the right. The system is parameterized by the position x of the cart, the angle of the pole ψ, the velocity $\dot{x}$ and the angular velocity $\dot{\psi}$.

experiment is given below.

Gravitational acceleration (g)	$9.8\frac{m}{s^2}$
Mass of the pole (m)	$0.5kg$
Mass of the cart (M)	$0.5kg$
Length of the pole (l)	$0.6m$
Friction coefficient (b)	$0.1N(ms)^{-1}$
Integration time step (Δt)	$0.1s$
Standard deviation of z (σ_2)	0.01
Discount factor (γ)	0.95

α_t	$t^{-1.0}$
β_t	$t^{-0.6}$
c_t	0.01
ϵ_1	0.95

The results of the experiments are shown in Figure 5.

4.2. Experiment 2: 5-Link Actuated Pendulum Balancing [14]

Setup: 5 independent poles each with mass m and length l with the top pole being a pendulum connected using 5 rotational joints.

Goal: To keep all the poles in the upright position by applying independent torques at each joint.

State space: The state $s = (q, \dot{q})^\top \in \mathbb{R}^{10}$ where $q = (\psi_1, \psi_2, \psi_3, \psi_4, \psi_5) \in \mathbb{R}^5$ and $\dot{q} = (\dot{\psi}_1, \dot{\psi}_2, \dot{\psi}_3, \dot{\psi}_4, \dot{\psi}_5) \in \mathbb{R}^5$ where ψ_i is the angle of the pole i *w.r.t.* the vertical axis and $\dot{\psi}_i$ is the angular velocity.

Control space: The action $a = (a_1, a_2, \ldots, a_5)^\top \in \mathbb{R}^5$ where a_i is the torque applied to the joint i. The

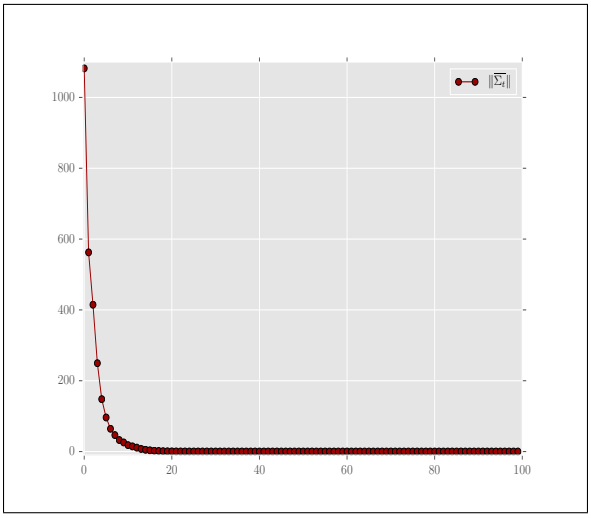

(a) $\|\bar{\Sigma}_t\|_F$ (where $\|\cdot\|_F$ is the Frobenius norm)

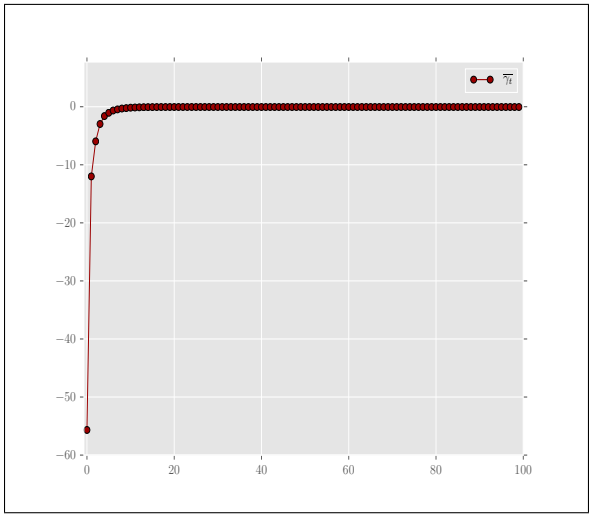

(b) $\bar{\gamma}_t^*$

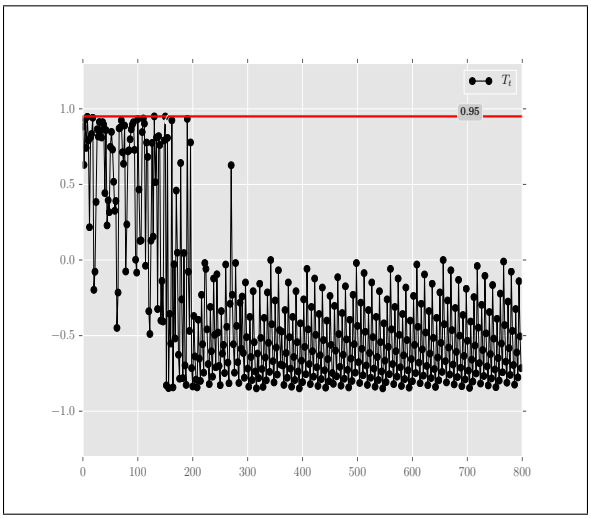

(c) T_t

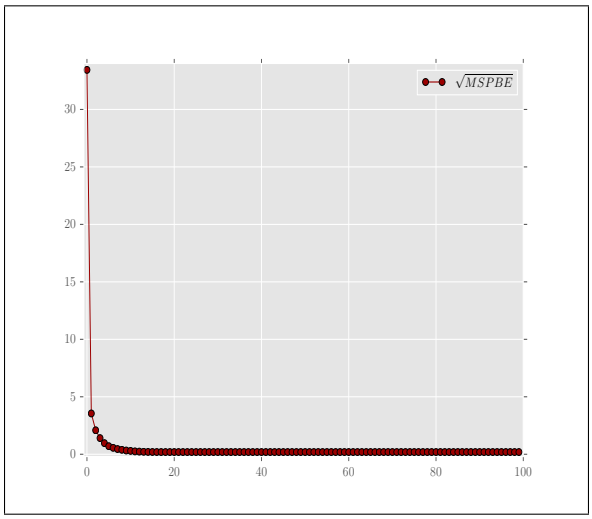

(d) $\sqrt{\mathrm{MSPBE}(\bar{\mu}_t)}$

Figure 5 The Cart-Pole setting. The evolutionary trajectory of the variables $\|\bar{\Sigma}_t\|_F$ (where $\|\cdot\|_F$ is the Frobenius norm), $\bar{\gamma}_t^*$, T_t and $\sqrt{\mathrm{MSPBE}(\bar{\mu}_t)}$. Note that both $\bar{\gamma}_t^*$ and $\sqrt{\mathrm{MSPBE}(\bar{\mu}_t)}$ converge to 0 as $t \to \infty$, while $\|\bar{\Sigma}_t\|_F$ also converges to 0. This implies that the model $\bar{\theta}_t = (\bar{\mu}_t, \bar{\Sigma}_t)^\top$ converges to the degenerate distribution concentrated on z^*. The evolutionary track of T_t shows that T_t does not cross the $\epsilon_1 = 0.95$ line after the model $\bar{\theta}_t = (\bar{\mu}_t, \bar{\Sigma}_t)^\top$ reaches a close neighbourhood of its limit.

stochastic policy used in this setting corresponds to $\pi(a|s) = \mathcal{N}(a|\beta_1^\top s, \sigma_1^2)$.

System dynamics: The approximate linear system dynamics is given by

$$\begin{bmatrix} q_{t+1} \\ \dot{q}_{t+1} \end{bmatrix} = \begin{bmatrix} I & \Delta t\, I \\ -\Delta t\, M^{-1}U & I \end{bmatrix} \begin{bmatrix} q_t \\ \dot{q}_t \end{bmatrix} + \Delta t \begin{bmatrix} 0 \\ M^{-1} \end{bmatrix} a + \mathbf{z} \tag{91}$$

where Δt is the integration time step, *i.e.*, the time difference between two transitions, M is the mass matrix in the upright position where $M_{ij} = l^2(6 - max(i,j))m$ and U is a diagonal matrix with $U_{ii} = -gl(6-i)m$. Each component of $\mathbf{z}$ is a Gaussian noise.

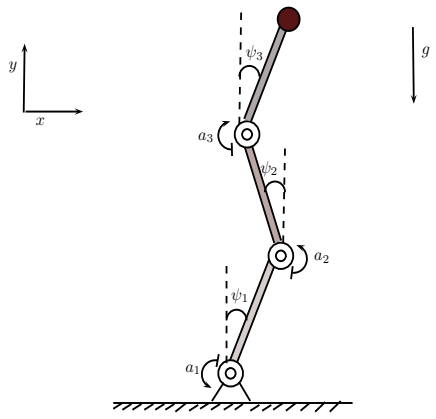

Figure 6 3-link actuated pendulum setting. Each rotational joint i, $1 \leq i \leq 3$ is actuated by a torque a_i . The system is parameterized by the angle ψ_i against the vertical direction and the angular velocity $\dot{\psi}_i$. The goal is to balance the pole in the upright direction, *i.e.*, all ψ_i should be as close to 0 as possible.

Reward function: $\mathrm{R}(q, \dot{q}, a) = -q^\top q$.

Feature vectors: $\phi(s \in \mathbb{R}^{10}) = (1, s_1^2, s_2^2 \ldots, s_1 s_2, s_1 s_3, \ldots, s_9 s_{10})^\top \in \mathbb{R}^{46}$.

Evaluation policy: The policy evaluated in the experiment is the optimal policy $\pi^*(a|s) = \mathcal{N}(a|{\beta_1^*}^\top s, \sigma_1^{*2})$. The parameters β_1^* and σ_1^* are computed using dynamic programming. The feature set chosen above is a perfect feature set, *i.e.*, $V^{\pi^*} \in \{\Phi z | z \in \mathbb{R}^k\}$.

The table of the various parameter values we used in our experiment is given below. Note that we have used constant step-sizes in this experiment.

Gravitational acceleration (g)	$9.8\frac{m}{s^2}$
Mass of the pole (m)	$1.0kg$
Length of the pole (l)	$1.0m$
Integration time step (Δt)	$0.1s$
Discount factor (γ)	0.95

α_t	0.001
β_t	0.05
c_t	0.05
ϵ_1	0.95

The results of the experiment are shown in Figure 7.

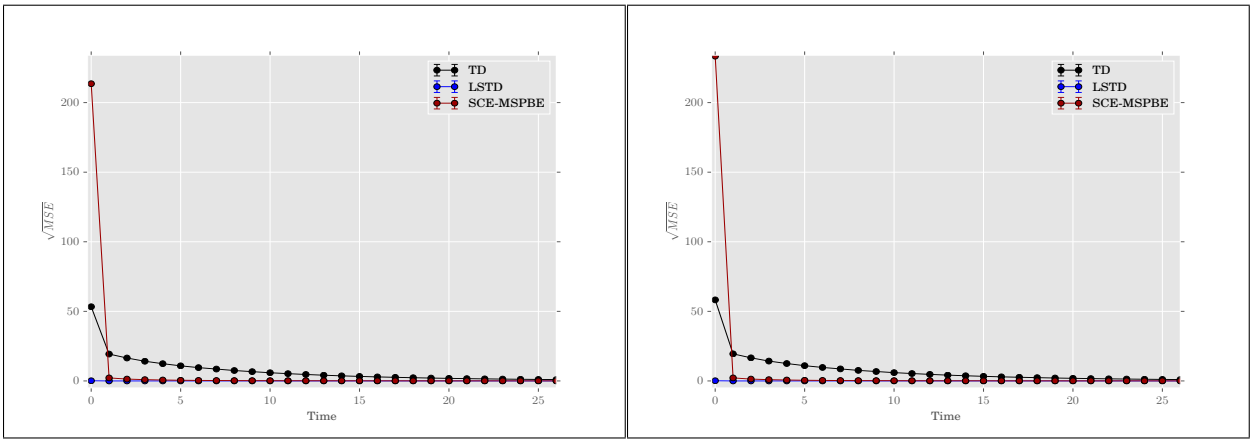

(a) $\sqrt{\text{MSPBE}(\bar{\mu}_t)}$ (b) $\sqrt{\text{MSE}(\bar{\mu}_t)}$

Figure 7 5-link actuated pendulum setting. The respective trajectories of the $\sqrt{\text{MSPBE}}$ and $\sqrt{\text{MSE}}$ generated by TD(0), LSTD(0) and SCE-MSPBEM algorithms are plotted. The graph on the left is for $\sqrt{\text{MSPBE}}$, while on the right is that of $\sqrt{\text{MSE}}$. Note that $\sqrt{\text{MSE}}$ also converges to 0 since the feature set is perfect.

4.3. Experiment 3: Baird's 7-Star MDP [10]

Our algorithm was also tested on Baird's star problem [10] with $|\mathbb{S}| = 7$, $|\mathbb{A}| = 2$ and $k = 8$. We let ν be the uniform distribution over $\mathbb{S}$ and the feature matrix Φ and the transition matrix P^π are given by

$$\Phi = \begin{pmatrix} 1\ 2\ 0\ 0\ 0\ 0\ 0\ 0 \\ 1\ 0\ 2\ 0\ 0\ 0\ 0\ 0 \\ 1\ 0\ 0\ 2\ 0\ 0\ 0\ 0 \\ 1\ 0\ 0\ 0\ 2\ 0\ 0\ 0 \\ 1\ 0\ 0\ 0\ 0\ 2\ 0\ 0 \\ 1\ 0\ 0\ 0\ 0\ 0\ 2\ 0 \\ 2\ 0\ 0\ 0\ 0\ 0\ 0\ 1 \end{pmatrix} \qquad \mathrm{P}^{\pi} = \begin{pmatrix} 0\ 0\ 0\ 0\ 0\ 0\ 1 \\ 0\ 0\ 0\ 0\ 0\ 0\ 1 \\ 0\ 0\ 0\ 0\ 0\ 0\ 1 \\ 0\ 0\ 0\ 0\ 0\ 0\ 1 \\ 0\ 0\ 0\ 0\ 0\ 0\ 1 \\ 0\ 0\ 0\ 0\ 0\ 0\ 1 \\ 0\ 0\ 0\ 0\ 0\ 0\ 1 \end{pmatrix}.$$

The reward function is given by $\mathrm{R}(s, s') = 0$, $\forall s, s' \in \mathbb{S}$. The Markov Chain in this case is not ergodic and hence belongs to an off-policy setting. This is a classic example where TD(0) is seen to diverge [10]. The performance comparison of the algorithms GTD2, TD(0) and LSTD(0) with SCE-MSPBEM is shown in Figure 9. The performance metric used here is the $\sqrt{MSE(\cdot)}$ of the prediction vector returned by the corresponding algorithm at time t. The algorithm parameters for the problem are given below:

α_t	0.001
β_t	0.05
c_t	0.01
ϵ_1	0.8

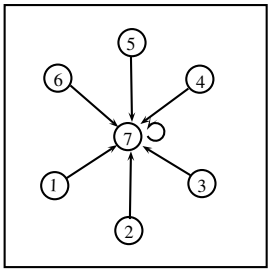

Figure 8: Baird's 7-star MDP

A careful analysis in [34] has shown that when the discount factor $\gamma \leq 0.88$, with appropriate learning rate, TD(0) converges. Nonetheless, it is also shown in the same paper that for discount factor $\gamma = 0.9$, TD(0) will diverge for all values of the learning rate. This is explicitly demonstrated in Figure 9. However our algorithm SCE-MSPBEM converges in both cases, which demonstrates the stable behaviour exhibited by our algorithm.

The algorithms were also compared on the same Baird's 7-star, but with a different feature matrix Φ_1.

$$\Phi_1 = \begin{pmatrix} 1 & 2 & 0 & 0 & 0 & 0 & 1 & 0 \\ 1 & 0 & 2 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 2 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 2 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 2 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 3 \\ 2 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix}.$$

In this case, the reward function is given by $\mathrm{R}(s, s') = 2.0$, $\forall s, s' \in \mathbb{S}$. Note that Φ_1 gives an imperfect feature set. The algorithm parameter values used are same as earlier. The results are show in Figure 10. In this case also, TD(0) diverges. However, SCE-MSPBEM is exhibiting good stable behaviour.

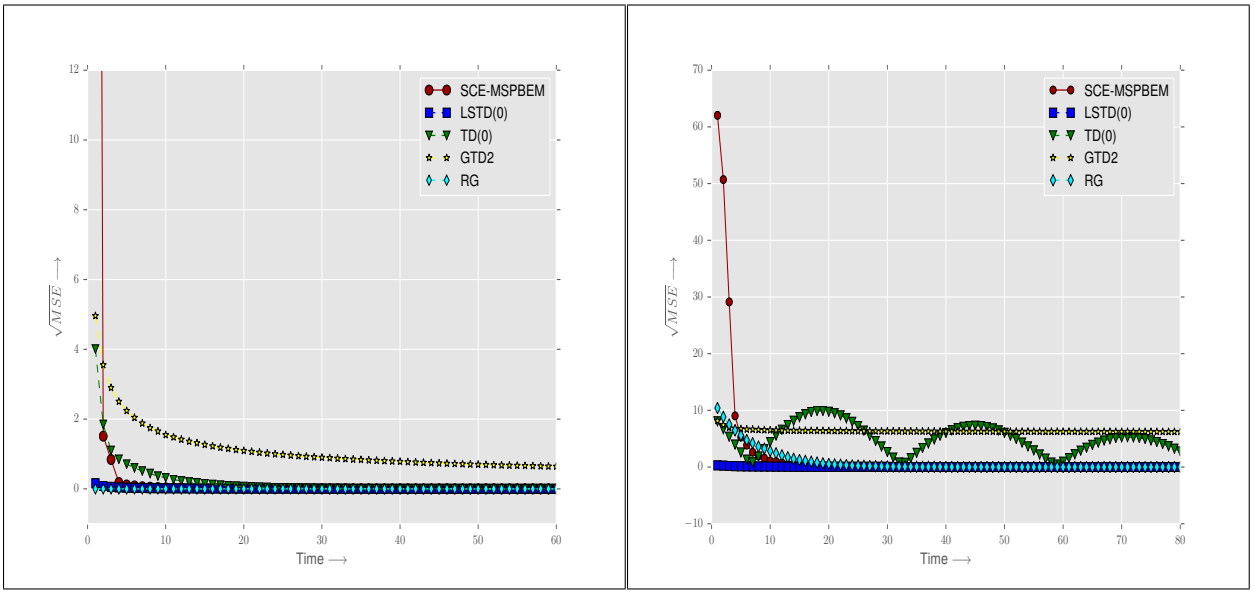

(a) Discount factor $\gamma = 0.1$ (b) Discount factor $\gamma = 0.9$

Figure 9 Baird's 7-Star MDP with perfect feature set. For $\gamma = 0.1$, all the algorithms show almost the same rate of convergence. The initial jump of SCE-MSPBEM is due to the fact that the initial value is far from the limit. For $\gamma = 0.9$, TD(0) does not converge and GTD2 is slower. However, SCE-MSPBEM exhibits good convergence behaviour.

Experiment 4: 10-State Ring MDP [33]

Next, we studied the performance comparisons of the algorithms on a 10-ring MDP with $|\mathbb{S}| = 10$ and $k = 8$. We let ν be the uniform distribution over $\mathbb{S}$. The transition matrix P^π and the feature matrix Φ are

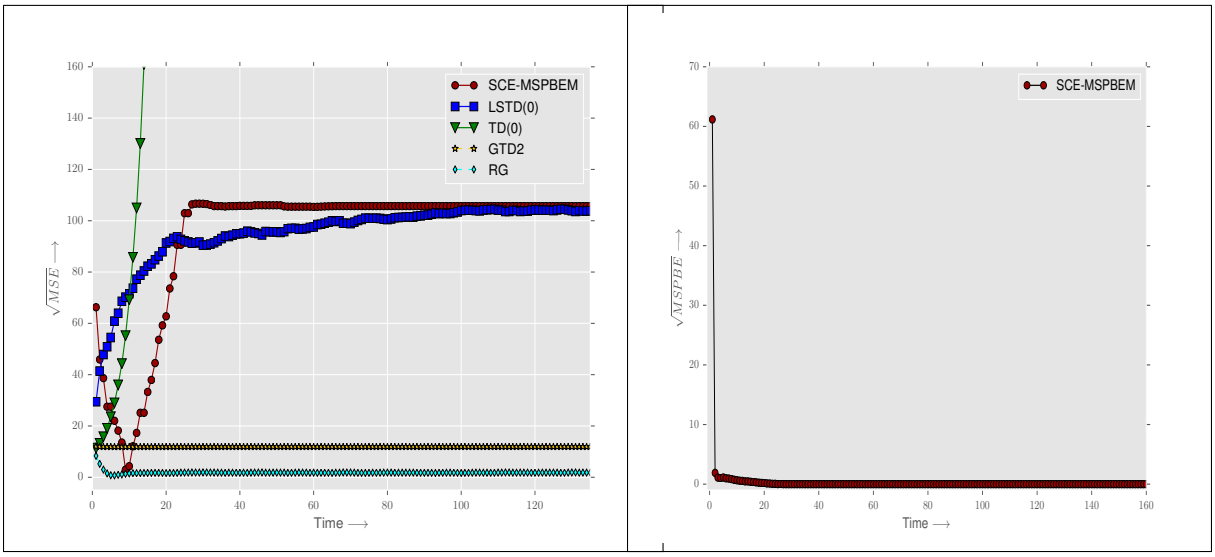

(a) $\sqrt{\text{MSE}}$. (b) $\sqrt{\text{MSPBE}}$.

Figure 10 Baird's 7-Star MDP with imperfect feature set. Here the discount factor $\gamma = 0.99$. In this case, TD(0) method diverges. However, $\sqrt{\text{MSE}}$ of SCE-MSPBEM and LSTD(0) converge to the same limit point (=103.0), while RG converges to a different limit (= 1.6919). This is because the feature set is imperfect and also due to the fact that RG minimizes MSBR, while SCE-MSPBEM and LSTD minimize MSPBE. To verify this fact, note that in (b), $\sqrt{\text{MSPBE}(\bar{\mu}_t)}$ of SCE-MSPBEM converges to 0.

given by

$$\mathrm{P}^{\pi} = \begin{pmatrix} 0&1&0&0&0&0&0&0&0&0\\ 0&0&1&0&0&0&0&0&0&0\\ 0&0&0&1&0&0&0&0&0&0\\ 0&0&0&0&1&0&0&0&0&0\\ 0&0&0&0&0&1&0&0&0&0\\ 0&0&0&0&0&0&1&0&0&0\\ 0&0&0&0&0&0&0&1&0&0\\ 0&0&0&0&0&0&0&0&1&0\\ 0&0&0&0&0&0&0&0&0&1\\ 1&0&0&0&0&0&0&0&0&0 \end{pmatrix}, \quad \Phi = \begin{pmatrix} 1&0&0&0&0&0&0&0\\ 0&1&0&0&0&0&0&0\\ 0&0&1&0&0&0&0&0\\ 0&0&0&1&0&0&0&0\\ 0&0&0&0&1&0&0&0\\ 0&0&0&0&0&1&0&0\\ 0&0&0&0&0&0&1&0\\ 0&0&0&0&0&0&0&1\\ 0&0&0&0&0&0&0&1\\ 0&0&0&0&0&1&0&0 \end{pmatrix}.$$

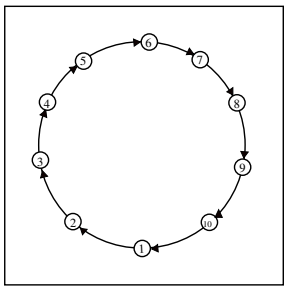

Figure 11: 10-Ring MDP

The reward function is $\mathrm{R}(s, s') = 1.0, \forall s, s' \in \mathbb{S}$. The performance comparisons of the algorithms GTD2, TD(0) and LSTD(0) with SCE-MSPBEM are shown in Figure 12. The performance metric used here is the $\sqrt{MSE(\cdot)}$ of the prediction vector returned by the corresponding algorithm at time t.

The algorithm parameters for the problem are as follows:

α_t	0.001
β_t	0.05
c_t	0.075
ϵ_1	0.85

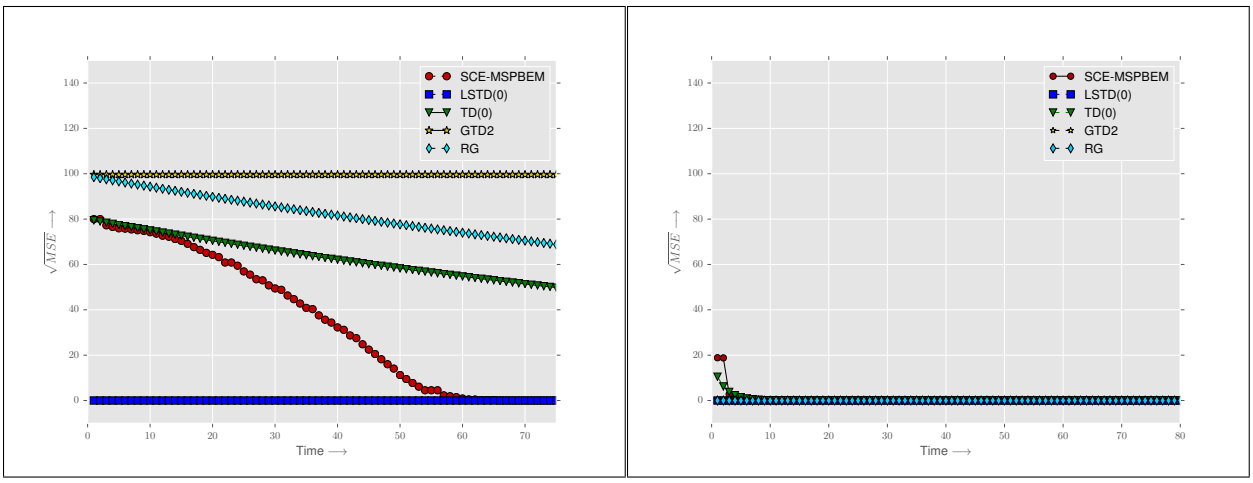

(a) Discount factor $\gamma = 0.99$ (b) Discount factor $\gamma = 0.1$

Figure 12 10-Ring MDP with perfect feature set: For $\gamma = 0.1$, all the algorithms exhibit almost the same rate of convergence. For $\gamma = 0.99$, SCE-MSPBEM converges faster than TD(0), GTD2 and RG.

4.4. Experiment 5: Large State Space Random MDP with Radial Basis Functions and Fourier Basis

These experiments were designed by us. Here, tests were performed by varying the feature set to prove that the algorithm is not dependent on any particular feature set. Two types of feature sets are used here: Fourier Basis Functions and Radial Basis Functions (RBF).

Figure 13 shows the performance comparisons when Fourier Basis Functions are used for the features $\{\phi_i\}_{i=1}^{k}$, where

$$\phi_i(s) = \begin{cases} 1 & \text{if } i = 1, \\ \cos \frac{(i+1)\pi s}{2} & \text{if } i \text{ is odd}, \\ \sin \frac{i\pi s}{2} & \text{if } i \text{ is even}. \end{cases} \tag{92}$$

Figure 14 shows the performance comparisons when RBF is used instead for the features $\{\phi_i\}_{i=1}^k$, where

$$\phi_i(s) = e^{-\frac{(s-m_i)^2}{2.0v_i^2}} \tag{93}$$

with m_i and v_i fixed *a priori*.

In both the cases, the reward function is given by

$$\mathrm{R}(s,s') = G(s)G(s')\left(\frac{1}{(1.0+s\prime)^{0.25}}\right), \qquad \forall s,s' \in \mathbb{S}, \tag{94}$$

where the vector $G \in (0,1)^{|\mathbb{S}|}$ is initialized for the algorithm with $G(s) \sim U(0,1), \forall s \in \mathbb{S}$.

Also in both the cases, the transition probability matrix P^π is generated as follows

$$\mathrm{P}^\pi(s,s') = \binom{|\mathbb{S}|}{s'} b(s)^{s'}(1.0-b(s))^{|\mathbb{S}|-s'}, \qquad \forall s,s' \in \mathbb{S}, \tag{95}$$

where the vector $b \in (0,1)^{|\mathbb{S}|}$ is initialized for the algorithm with $b(s) \sim U(0,1), \forall s \in \mathbb{S}$. It is easy to verify that the Markov Chain defined by P^π is ergodic in nature.

In the case of RBF, we have set $|\mathbb{S}| = 1000$, $|\mathbb{A}| = 200$, $k = 50$, $m_i = 10 + 20(i-1)$ and $v_i = 10$, while for Fourier Basis Functions, $|\mathbb{S}| = 1000$, $|\mathbb{A}| = 200$, $k = 50$. In both the cases, the distribution ν is the stationary distribution of the Markov Chain. The simulation is run sufficiently long to ensure that the chain achieves its steady state behaviour, *i.e.*, the states appear with the stationary distribution.

The algorithm parameters for the problem are as follows:

Both RBF & Fourier Basis	
α_t	0.001
β_t	0.05
c_t	0.075
ϵ_1	0.85

Also note that when Fourier basis is used, the discount factor $\gamma = 0.9$ and for RBFs, $\gamma = 0.01$. SCE-MSPBEM exhibits good convergence behaviour in both cases, which shows the non-dependence of SCE-MSPBEM on the discount factor γ. This is important because in [35], the performance of TD methods is shown to be dependent on the discount factor γ.

To get a measure of how well our algorithm performs on much larger problems, we applied it on a large MDP where $|\mathbb{S}| = 2^{15}$, $|\mathbb{A}| = 50$, $k = 100$ and $\gamma = 0.9$. The reward function R and the transition probability

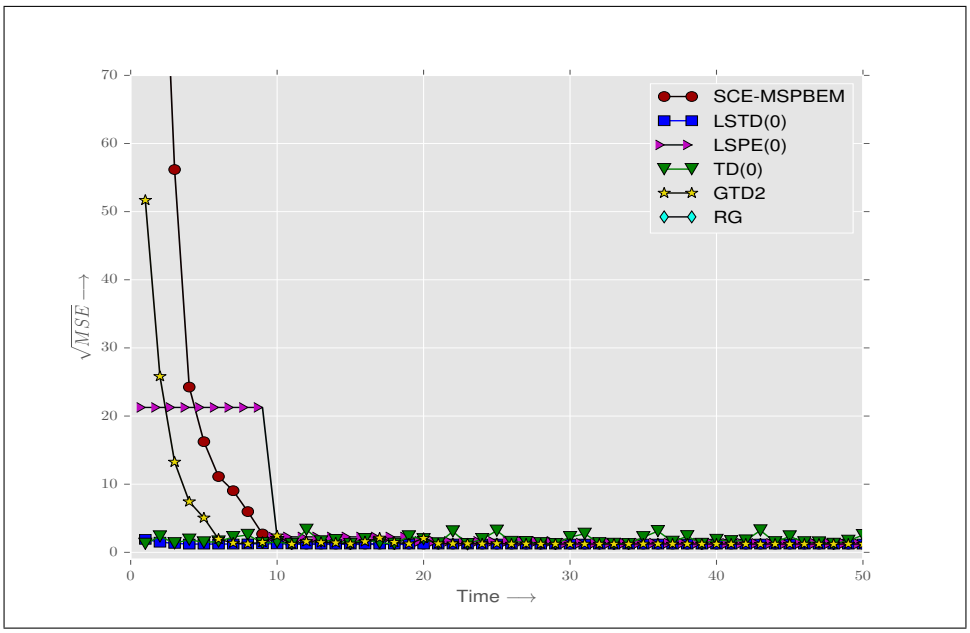

Figure 13 Fourier Basis Function: Here, $|\mathbb{S}| = 1000$, $|\mathbb{A}| = 200$, $k = 50$ and $\gamma = 0.9$. In this case, SCE-MSPBEM shows good convergence behaviour.

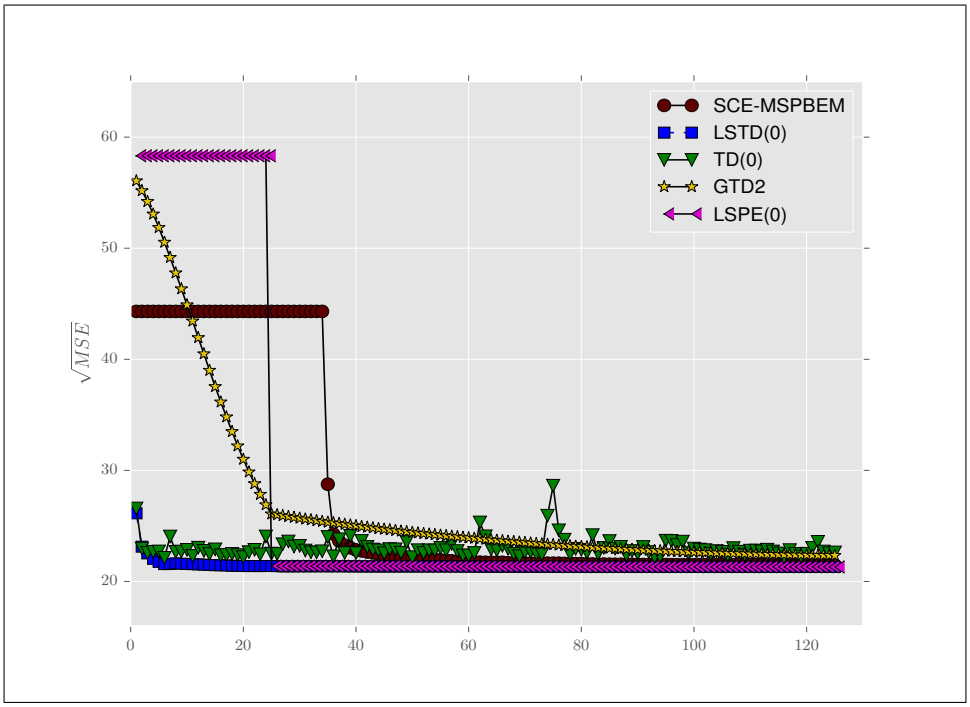

Figure 14 Radial Basis Function. Here, $|\mathbb{S}| = 1000$, $|\mathbb{A}| = 200$, $k = 50$ and $\gamma = 0.01$. In this case, SCE-MSPBEM converges to the same limit point as other algorithms.

matrix P^{π} are generated using the equations (94) and (95) respectively. RBFs are used as the features in this case. Since the MDP is huge, the algorithms were run on Amazon cloud servers. The true value function V^{π} was computed and the $\sqrt{MSE}$s of the prediction vectors generated by the different algorithms were compared. The performance results are shown in Table 2.

Ex#	SCE-MSPBEM	LSTD(0)	TD(0)	RLSTD(0)	LSPE(0)	GTD2
1	23.3393204	23.3393204	24.581849219	23.3393204	23.354929410	24.93208571
2	23.1428622	23.1428622	24.372033722	23.1428622	23.178814326	24.75593565
3	23.3327844	23.3327848	24.537556372	23.3327848	23.446585398	24.88119648
4	22.9786909	22.9786909	24.194543862	22.9786909	22.987520761	24.53206023
5	22.9502660	22.9502660	24.203561613	22.9502660	22.965571900	24.55473382
6	23.0609354	23.0609354	24.253239213	23.0609354	23.084399716	24.60783237
7	23.2280270	23.2280270	24.481937450	23.2280270	23.244345617	24.83529005

Table 2 Performance comparison of various algorithms with large state space. Here $|\mathbb{S}| = 2^{15}, |\mathbb{A}| = 50, k = 100$, and $\gamma = 0.9$. RBF is used as the feature set. The feature set is imperfect. The entries in the table correspond to the $\sqrt{\mathrm{MSE}}$ values obtained from the respective algorithms on 7 different random MDPs. While the entries of SCE-MSPBEM, LSTD(0) and RLSTD(0) appear to be similar, they actually differed in decimal digits that are not shown here for lack of space.

5. Conclusion and Future Work

We proposed, for the first time, an application of the Cross Entropy (CE) method to the problem of prediction in Reinforcement Learning (RL) under the linear function approximation architecture. This task is accomplished by remodelling the original CE algorithm as a multi-timescale stochastic approximation algorithm and using it to minimize the Mean Squared Projected Bellman Error (MSPBE). The proof of convergence to the optimum value using the ODE method is also provided. The theoretical analysis is supplemented by extensive experimental evaluation which is shown to corroborate the claim. Experimental comparisons with the state-of-the-art algorithms show the superiority in the accuracy of our algorithm while being competitive enough with regard to computational efficiency and rate of convergence.

The algorithm can be extended to non-linear approximation settings also. In [36], a variant of the TD(0) algorithm is developed and applied in the non-linear function approximation setting, where the convergence to the local optima is proven. But we believe our approach can converge to the global optimum in the non-linear case because of the application of a CE-based approach, thus providing a better approximation to the value function. The algorithm can also be extended to the off-policy case [8, 9], where the sample trajectories are developed using a behaviour policy which is different from the target policy whose value function is approximated. This can be achieved by appropriately integrating a weighting ratio [37] in the recursions. TD learning methods are shown to be divergent in the off-policy setting [10]. So it will be interesting to see how our algorithm behaves in such a setting. Another future work includes extending this optimization technique to the control problem to obtain an optimum policy. This can be achieved by parametrizing the policy space and using the optimization technique in SCE-MSPBEM to search in this parameter space.

References

[1] R. S. Sutton and A. G. Barto, *Introduction to reinforcement learning*. MIT Press, New York, USA, 1998.

[2] D. J. White, "A survey of applications of Markov decision processes," *Journal of the Operational Research Society*, pp. 1073–1096, 1993.

[3] D. P. Bertsekas, *Dynamic programming and optimal control*, vol. 2. Athena Scientific Belmont, USA, 2013.

[4] M. G. Lagoudakis and R. Parr, "Least-squares policy iteration," *The Journal of Machine Learning Research*, vol. 4, pp. 1107–1149, 2003.

[5] G. Konidaris, S. Osentoski, and P. S. Thomas, "Value Function Approximation in Reinforcement Learning Using the Fourier Basis.," in *AAAI*, 2011.

[6] M. Eldracher, A. Staller, and R. Pompl, *Function approximation with continuous valued activation functions in CMAC*. Citeseer, 1994.

[7] R. S. Sutton, "Learning to predict by the methods of temporal differences," *Machine learning*, vol. 3, no. 1, pp. 9–44, 1988.

[8] R. S. Sutton, H. R. Maei, and C. Szepesvári, "A convergent $o(n)$ temporal-difference algorithm for off-policy learning with linear function approximation," in *Advances in neural information processing systems*, pp. 1609–1616, 2009.

[9] R. S. Sutton, H. R. Maei, D. Precup, S. Bhatnagar, D. Silver, C. Szepesvári, and E. Wiewiora, "Fast gradient-descent methods for temporal-difference learning with linear function approximation," in *Proceedings of the 26th Annual International Conference on Machine Learning*, pp. 993–1000, ACM, 2009.

[10] L. Baird, "Residual algorithms: Reinforcement learning with function approximation," in *Proceedings of the twelfth international conference on machine learning*, pp. 30–37, 1995.

[11] S. J. Bradtke and A. G. Barto, "Linear least-squares algorithms for temporal difference learning," *Machine Learning*, vol. 22, no. 1-3, pp. 33–57, 1996.

[12] J. A. Boyan, "Technical update: Least-squares temporal difference learning," *Machine Learning*, vol. 49, no. 2-3, pp. 233–246, 2002.

[13] A. Nedić and D. P. Bertsekas, "Least squares policy evaluation algorithms with linear function approximation," *Discrete Event Dynamic Systems*, vol. 13, no. 1-2, pp. 79–110, 2003.

[14] C. Dann, G. Neumann, and J. Peters, "Policy evaluation with temporal differences: A survey and comparison," *The Journal of Machine Learning Research*, vol. 15, no. 1, pp. 809–883, 2014.

[15] J. N. Tsitsiklis and B. Van Roy, "An analysis of temporal-difference learning with function approximation," *Automatic Control, IEEE Transactions on*, vol. 42, no. 5, pp. 674–690, 1997.

[16] R. J. Williams and L. C. Baird, "Tight performance bounds on greedy policies based on imperfect value functions," tech. rep., Citeseer, 1993.

[17] B. Scherrer, "Should one compute the Temporal Difference fix point or minimize the Bellman Residual? the unified oblique projection view," in *27th International Conference on Machine Learning-ICML 2010*, 2010.

[18] M. Zlochin, M. Birattari, N. Meuleau, and M. Dorigo, "Model-based search for combinatorial optimization: A critical survey," *Annals of Operations Research*, vol. 131, no. 1-4, pp. 373–395, 2004.

[19] J. Hu, M. C. Fu, and S. I. Marcus, "A model reference adaptive search method for global optimization," *Operations Research*, vol. 55, no. 3, pp. 549–568, 2007.

[20] E. Zhou, S. Bhatnagar, and X. Chen, "Simulation optimization via gradient-based stochastic search," in *Simulation Conference (WSC), 2014 Winter*, pp. 3869–3879, IEEE, 2014.

[21] M. Dorigo and L. M. Gambardella, "Ant colony system: a cooperative learning approach to the traveling salesman problem," *Evolutionary Computation, IEEE Transactions on*, vol. 1, no. 1, pp. 53–66, 1997.

[22] H. Mühlenbein and G. Paass, "From recombination of genes to the estimation of distributions i. binary parameters," in *Parallel Problem Solving from NaturePPSN IV*, pp. 178–187, Springer, 1996.

[23] J. Hu, M. C. Fu, and S. I. Marcus, "A model reference adaptive search method for stochastic global optimization," *Communications in Information & Systems*, vol. 8, no. 3, pp. 245–276, 2008.

[24] I. Menache, S. Mannor, and N. Shimkin, "Basis function adaptation in temporal difference reinforcement learning," *Annals of Operations Research*, vol. 134, no. 1, pp. 215–238, 2005.

[25] R. Y. Rubinstein and D. P. Kroese, *The cross-entropy method: a unified approach to combinatorial optimization, Monte-Carlo simulation and machine learning*. Springer Science & Business Media, 2013.

[26] P.-T. De Boer, D. P. Kroese, S. Mannor, and R. Y. Rubinstein, "A tutorial on the cross-entropy method," *Annals of operations research*, vol. 134, no. 1, pp. 19–67, 2005.

[27] J. Hu and P. Hu, "On the performance of the cross-entropy method," in *Simulation Conference (WSC), Proceedings of the 2009 Winter*, pp. 459–468, IEEE, 2009.

[28] V. S. Borkar, "Stochastic approximation: A dynamical systems viewpoint," *Cambridge University Press*, 2008.

[29] H. J. Kushner and D. S. Clark, *Stochastic approximation for constrained and unconstrained systems*. Springer Verlag, New York, 1978.

[30] H. Robbins and S. Monro, "A stochastic approximation method," *The Annals of Mathematical Statistics*, pp. 400–407, 1951.

[31] T. Homem-de Mello, "A study on the cross-entropy method for rare-event probability estimation," *INFORMS Journal on Computing*, vol. 19, no. 3, pp. 381–394, 2007.

[32] C. Kubrusly and J. Gravier, "Stochastic approximation algorithms and applications," in *1973 IEEE Conference on Decision and Control including the 12th Symposium on Adaptive Processes*, no. 12, pp. 763–766, 1973.

[33] B. Kveton, M. Hauskrecht, and C. Guestrin, "Solving factored mdps with hybrid state and action variables.," *J. Artif. Intell. Res.(JAIR)*, vol. 27, pp. 153–201, 2006.

[34] R. Schoknecht and A. Merke, "Convergent combinations of reinforcement learning with linear function approximation," in *Advances in Neural Information Processing Systems*, pp. 1579–1586, 2002.

[35] R. Schoknecht and A. Merke, "TD(0) converges provably faster than the residual gradient algorithm," in *ICML*, pp. 680–687, 2003.

[36] H. R. Maei, C. Szepesvári, S. Bhatnagar, D. Precup, D. Silver, and R. S. Sutton, "Convergent temporal-difference learning with arbitrary smooth function approximation," in *Advances in Neural Information Processing Systems*, pp. 1204–1212, 2009.

[37] P. W. Glynn and D. L. Iglehart, "Importance sampling for stochastic simulations," *Management Science*, vol. 35, no. 11, pp. 1367–1392, 1989.